Ultra-compacte geïntegreerde optische filters in silicium-op-isolator op basis van waferschaaltechnologie

Ultra-Compact Integrated Optical Filters in Silicon-on-insulator by Means of Wafer-Scale Technology

Pieter Dumon

UNIVERSITEIT
GENT

Promotor:
Prof. dr. ir. R. Baets   Universiteit Gent, INTEC

Examencommissie:
Prof. dr. ir. D. Dezutter, voorzitter   Universiteit Gent, INTEC
Prof. dr. P.I. Borel   Technical University of Denmark, COM
Dr. P. Absil   IMEC
Prof. dr. ir. K. Neyts   Universiteit Gent, ELIS
Dr. ir. W. Bogaerts   Universiteit Gent, INTEC
Prof. Dr. ir. D. Van Thourhout   Universiteit Gent, INTEC
Prof. Dr. ir. P. Bienstman   Universiteit Gent, INTEC

Universiteit Gent
Faculteit Ingenieurswetenschappen

Vakgroep Informatietechnologie (INTEC)
Sint-Pietersnieuwstraat 41
9000 Gent
BELGIUM

Tel.: +32-9-264.33.19
Fax: +32-9-264.35.93

Voorwoord

Infrarood laserlicht, onzichtbaar voor ons oog, dat mooi microscopisch kleine baantjes volgt, en in een zeker mate doet wat je het wil laten doen, het is voor mij ook na vier jaar onderzoek nog steeds iets heel fascinerend. Daarom vooral ben ik dit beginnen doen en heb ik het steeds graag gedaan. Dit bestuderen geeft ook een andere interpretatie aan wat we *groot* en *klein* noemen: de lichtbaantjes, golfgeleiders, zijn een paar tienduizendsten van een millimeter breed, of meer dan 100 keer fijner dan een mensenhaar. De *golflengtefilters* die we ermee maken zijn een duizendste tot een tiende van een vierkante millimeter groot, de oppervlakte van een naaldpunt tot een kleine speldenkop dus. Om die filters goed te laten werken, moeten de golfgeleiders echter gemaakt worden met een nauwkeurigheid van een nanometer, een miljoenste van een millimeter. Dit verschil in grootteordes is eveneens fascinerend.

In dit werk worden de golfgeleiders gemaakt in een plakje silicium, het materiaal van de micro-elektronica-industrie, met machines en door mensen in het onderzoekcentrum voor micro-elektronica (IMEC) in Leuven. Het zijn de dure toestellen en de werkmethodes in zo een omgeving die het mogelijk maken dat we nu reeds een nauwkeurigheid van de golfgeleiders van 5 nm tot 10 nm halen over meer dan honderd plakjes (chips) die in één keer gemaakt worden op een siliciumschijf van 200 mm diameter. Dit is de kracht van de elektronica-industrie, en het is deze kracht die we net willen aanwenden voor fotonische chips.

Een doctoraatswerk is zeer vakspecialistisch: je werkt in een deel van een deel van een deeldomeintje van de techniek en wetenschap. Het leuke daaraan, althans in de fotonicawereld, is dat er weinig mensen ter wereld zijn die identiek hetzelfde trachten te doen. Maar toch is het zeer verrijkend om met collega's samen te werken die een andere invalshoek hebben. Dit kon ik doen in het project met de ESA.

Freggels, ze leerden me een beetje voetballen en de wereld het 'Dumonke' kennen — en dan nu de allerbelangrijkste mensen: **mama en papa, zus en broer**, voor alles — **mijn grootouders**, ik was er altijd welkom en kon er in het weekend even de rust terugvinden — **Leen**, om mij te vinden.

Pieter Dumon
Gent, 16 februari 2007

Samenvatting - Summary

Samenvatting

Door gebruik te maken van silicium-op-isolator technologie, kunnen ge-
integreerde optische filters spectaculair verkleind worden. De technologie
laat toe licht te doen buigen, splitsen en interageren over een afstand van
slechts een paar duizendsten van een millimeter. De reden hiervoor is het
zeer hoge brekingsindexcontrast van silicium-op-isolator (SOI) golfgelei-
der structuren. Deze golfgeleiders hebben dimensies van minder dan een
micrometer, en sluiten het licht zeer sterk op in hun kern.

Geïntegreerde optische filters vormen een belangrijk onderdeel van
optische netwerken. In de kern van telecommunicatienetwerken, waar
gegevens aan grote snelheid getransporteerd worden, wordt golflengte
multiplexering (WDM) gebruikt om de grote bandbreedte van optische
vezels te benutten. Optische filters splitsen en combineren golflengtekanalen
waar nodig. Maar ook in toegangsnetwerken tot telecommunicatiedien-
sten en in optische verbindingen over korte afstanden kunnen optische
filters een rol spelen. Door meerdere golflengtes te gebruiken om hoges-
nelheidssignalen te transporteren, kan ook daar een hogere totale band-
breedte en een grotere flexibiliteit behaald worden. In die laatste toepassin-
gen echter, zijn componenten met een hoge functionaliteit en compactheid
en een laag vermogenverbruik belangrijk.

De integratie van meer passieve functionaliteit op een enkele chip is
daarom een stuwende factor. De sterke reductie van de grootte van com-
ponenten in hoog indexcontrast SOI heeft echter een nadeel: de struc-
turen zijn zeer gevoelig voor kleine variaties in hun afmetingen, dit zelfs
voor afwijkingen van slechts één miljoenste van een millimeter. Deze
kleine variaties ontstaan door de imperfectie van de fabricagetechnolo-
gie, die slechts binnen bepaalde grenzen verbeterbaar is. De enige manier
waarop silicium-technologie in echte toepassingen zal kunnen gebruikt
worden, en tegelijkertijd de grootste drijfveer om silicium-technologie te
gebruiken, is te bouwen op de processen op waferschaal die ontwikkeld
zijn voor de fabricatie van elektronische chips, en meer specifiek CMOS
componenten. Integratie met CMOS op dezelfde drager wordt dan zelfs

mogelijk. Fotonische en elektronische geïntegreerde circuits (IC's) stellen echter zeer verschillende eisen aan de fabricagetechnologie, ten minste voor algemene hoog indexcontrast fotonische structuren. Met waferschaal CMOS processen kunnen grote aantallen componenten op korte tijd gefabriceerd worden. Maar die processen waren nooit bedoeld om in één enkele processtap structuren in het silicium te etsen die variëren van zeer geïsoleerde golfgeleiders tot dicht opeengepakte en periodische structuren. Voor fotonische geïntegreerde circuits is dit nodig om alignatieproblemen te vermijden. Dit heeft een grote impact op het ontwikkelen van fabricageprocessen voor fotonische IC's, en op de performantie van optische filters. Variaties op nanometerschaal van de structuurafmetingen in geïntegreerde filters geven aanleiding tot onjuiste beeldvorming en minder dan optimale interferentie. Dit resulteert in optisch vermogenverlies en overspraak tussen golflengtekanalen.

In dit werk werden geïntegreerde optische filters gerealiseerd met technologie op waferschaal, waaronder optische lithografie met een belichtingsgolflengte in het diep ultraviolet. De componenten zijn gebaseerd op fotonische draad golfgeleiders in SOI, die het licht opsluiten door middel van totale interne reflectie. Verschillende types filters werden geïmplementeerd om een goed zicht te krijgen op zowel de mogelijkheden als de zaken die voor verbetering vatbaar zijn. De bestudeerde filtertypes zijn goed gekend uit materiaalsystemen met een lager brekingsindexcontrast: ringresonatoren, filters gebaseerd op Mach-Zehnder interferometers zoals de zeer flexibele roosterfilters *(lattice filters)*, en gefaseerde golfgeleiderrij *(arrayed waveguide grating, AWG)* golflengterouteringscomponenten.

De performantie van de componenten werd bestudeerd op basis van een in huis ontwikkeld fabricageproces. Dit proces werd ontwikkeld om structuren te fabriceren met variërende periode en vulfactor, zowel als de meer geïsoleerde structuren als fotonische draad golfgeleiders. De golfgeleiderverliezen zijn typisch rond 3.5 dB/cm. Een ouder proces, dat echter onvoldoende in staat was periodieke structuren te fabriceren, gaf zelfs een verlies tot 2.4 dB/cm.

In dit werk werden ringresonatoren gedemonstreerd met een bochtstraal van slechts 1 tot 5 µm. We vervaardigden filters die één kanaal selecteren uit een rij, met een grote periodiciteit van het transmissiespectrum, lage verliezen en tot 20 dB extinctie in het doorgangskanaal. Met lagere verliezen en een lagere gedistribueerde koppeling tussen de modes in de ringresonator, zouden hogere orde filters ontworpen en gefabriceerd kunnen worden.
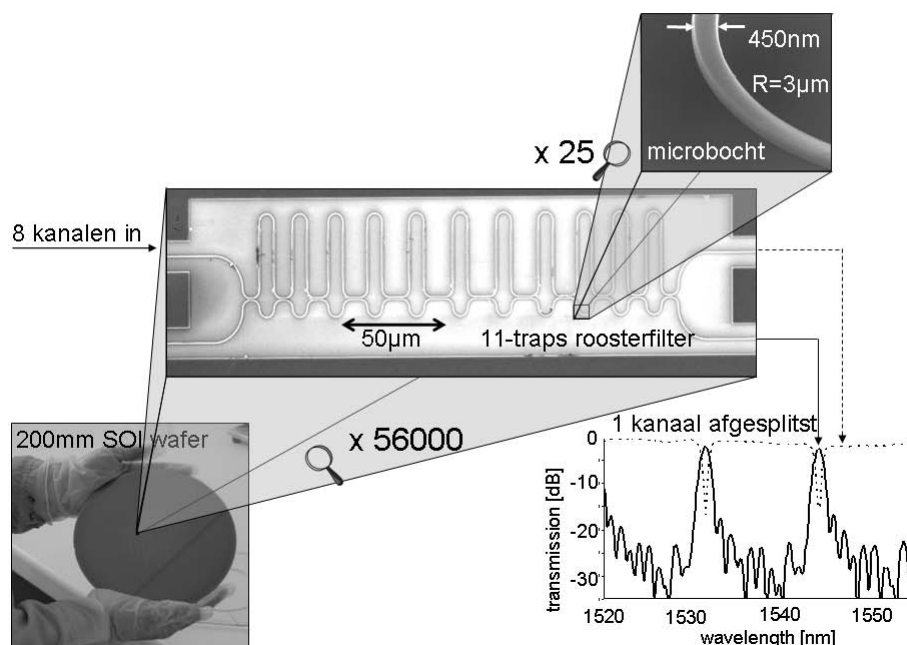
We ontwierpen en fabriceerden Mach-Zehnder roosterfilters die één kanaal selecteren uit 4 of 8 kanalen. Deze filters hadded WDM kanaalspatieringen van 100 tot 400 GHz. Ze kunnen ontworpen worden door middel van filtersynthesemethodes die wel bekend zijn uit de wereld van de digitale filters. De gefabriceerde componenten hadden een laag verlies, bandbreedtes zoals ontworpen, en een extinctie van de orde van 10 dB (tot 15 dB in de uitgang met het geselecteerde kanaal). Analyse van de karakterisatie toont dat de overspraak zowel te wijten is aan variaties van de koppelaars, als aan fasefouten in de vertragingslijnen in die filters. Deze fasefouten werden nochtans reeds gereduceerd door de golfgeleiders breder te maken in de rechte delen van de vertragingslijnen.

Golflengterouteringscomponenten gebaseerd op gefaseerde golfgeleiderrij *(arrayed waveguide grating)* structuren werden ontworpen en gefabriceerd. Componenten met 4 tot 16 kanalen, met kanaalspatiëringen van 100 tot 400 GHz, hadden verliezen tot 2.5 dB en isolatie tussen de kanalen tot 15 a 17 dB voor de beste componenten. Om dit te bereiken, werd lokaal het brekingsindexcontrast verlaagd rond de sterkoppelaar in die componenten. Hiervoor werd een etsproces met twee etsdieptes gebruikt. De golfgeleiderrij werd terug minder gevoelig gemaakt voor kleine afwijkingen van de golfgeleiderafmetingen door de golfgeleiders te verbreden.

We karakteriseerden de uniformiteit tussen identieke filters, gespatieerd over een afstand van een paar tientallen micrometer, een paar millimeter, en over verschillende chips op de wafer. Het resultaat wijst erop dat de stochastische variaties op korte lengteschaal binnen de golfgeleiders en tussen de golfgeleiders enige deterministische component overschaduwen. De uniformiteit op korte afstand is wel beduidend beter dan de uniformiteit over langere afstanden. Gecombineerd met de gemeten overspraakniveau's en de karakterisatie van het fabricageproces, de wafers en de lithografische maskers, wijst dit op breedtevariaties in de golfgeleiders op een lengteschaal van enkele honderden nanometer tot enkele micrometer. Deze zouden geëlimineerd moeten worden. Er zijn echter aanwijzingen dat deze variaties reeds gedeeltelijk aanwezig zijn op het lithografisch masker.

Naast het verbeteren van de performantie van de componenten, zijn de belangrijkste obstakels voor echte toepassingen het verbinden van de in- en uitgangen van de component naar optische vezels en de polarisatie-afhankelijkheid. Optische vezels en submicrometer golfgeleiders koppelen lijkt op het eerste zicht moeilijk. Echter, de voorbije jaren werden koppelingsmethodes met hoge performantie en goede aligneringstoleranties aangetoond. Deze werden net mogelijk, ietwat contra-intuïtief misschien, door het hoge brekingsindexcontrast. Het werkelijk fysisch verbinden van de siliciumchip met de vezels op een stabiele, reproduceerbare manier die vele parallelle vezels ondersteunt, vormt echter nog steeds een obstakel. Bovendien zou de vezelverbinding niet veel groter mogen zijn dan de eigenlijke component. In dit werk werd een AWG-gebaseerde golflengterouterings component verbonden met een standaard koppelstructuur met een rij van 8 vezels, op basis van verticale koppeling. Deze component is vervolgens gebruikt in een demonstratie-opstelling voor een hogesnelheids interconnectie toepassing. Deze toepassing was een herconfigureerbaar optisch netwerk voor de verbinding van meerdere elektronische borden aan 10 Gbps. Met de SOI golflengterouterings component als het hart van het netwerk, toonde de karakterisatie van het netwerk een schaalbaarheid van 4 to 9 poorten. Dit toont tegelijkertijd de limieten en de mogelijkheden van SOI componenten met huidige technologie aan. Dit optisch netwerk werd onderzocht in het kader van de verbinding van elektronische borden in satellieten in de ruimte. Daarom werd ook een eerste test gedaan naar hoe goed SOI componenten bestand zijn tegen hoogenergetische straling. Deze straling kan voor een verplaatsing van het transmissiespectrum zorgen. De tests wijzen erop dat alhoewel passieve geïntegreerde optische filters in SOI een hogere spectrale verplaatsing ver-

tonen dan hun broertjes in laag indexcontrast, de veranderingen slechts belangrijk worden op missies van lange duur of in een milieu met veel inval van straling. Bovendien kunnen de waargenomen veranderingen gecompenseerd worden door een aanpassing van de temperatuur van de chip.

De componenten werden gefabriceerd met processen op waferschaal. Deze processen zijn ontworpen voor de fabricage van een grote range aan optische structuren, en niet specifiek gericht op de bestudeerde (veelal geïsoleerde) structuren. Ondanks de huidige niet-uniformiteit tussen en binnenin de structuren, hebben we een performantie aangetoond die degelijk genoeg was voor eerste tests van toepassingen. De geïntegreerde optische filters hebben bovendien kanaalspatiëringen die compatibel zijn met WDM en beslaan slechts een chipoppervlak van 0.01 tot 0.1 mm^2.

Betrouwbare volumefabricage van fotonische geïntegreerde circuits door middel van waferschaal technologieën, met een hoge functionaliteit en met een grote variëteit aan onderliggende fysische structuren, is nog niet voor morgen. Maar we hebben toch een klein stukje van de af te leggen weg geplaveid, door ultra-compacte geïntegreerde optische filters in silicium-op-isolator aan te tonen, gebaseerd op fotonische draad golfgeleiders en gefabriceerd met waferschaaltechnologie, dit met een voldoende performantie voor eerste tests en als basis voor verdere ontwikkeling.

Summary

Silicon-on-insulator technology allows for a dramatic reduction of the footprint of integrated optical filters. Light can be forced to turn, split and interact over distances of just a few thousandths of a millimetre. The key to this size reduction is the very high refractive index contrast of silicon-on-insulator (SOI) waveguide structures. These waveguides have sub-micrometre dimensions and confine light very tightly within the waveguide core.

Integrated optical filters are one of the key components in optical networks. In the backplanes of telecommunication networks, wavelength division multiplexing (WDM) is used to exploit the large bandwidth of optical fibres. Optical filters are needed to split and combine wavelength channels. Also in access networks and in optical interconnects higher aggregate bandwidths and a higher flexibility could be obtained by using multiple wavelengths to transport high data-rate signals. However, for usage in these environments, components with a higher functionality, small footprint and lower power consumption are needed.

Integrating much more passive functionality on a single chip is therefore an enabling factor. The size reduction of high index-contrast SOI technology however comes at the cost of a very strong sensitivity to small dimensional variations, which exist due to processing imperfections. Also, the only way to use silicon technology in the real world, and at the same time the main reason to use it, is to build on the wafer-scale technology developed for fabrication of electronic device, and more specifically CMOS devices. Integration with CMOS on the same chip is then even a possibility. However, photonic and electronic integrated circuits have very different requirements with respect to fabrication technology, at least for general high-contrast integrated structures. The wafer-scale, high throughput CMOS technology was not developed to support features varying from fully isolated to dense and perfectly periodic structures in the same processing step. Due to the alignment tolerances of photonic circuits, this is needed for photonic ICs. This has a large impact on the devel-

opment of fabrication processes and the performance of optical filters. Small, nanometre-scale variations of the feature size within a structure give rise to imperfect imaging or suboptimal interference, resulting in loss and crosstalk between wavelength channels.
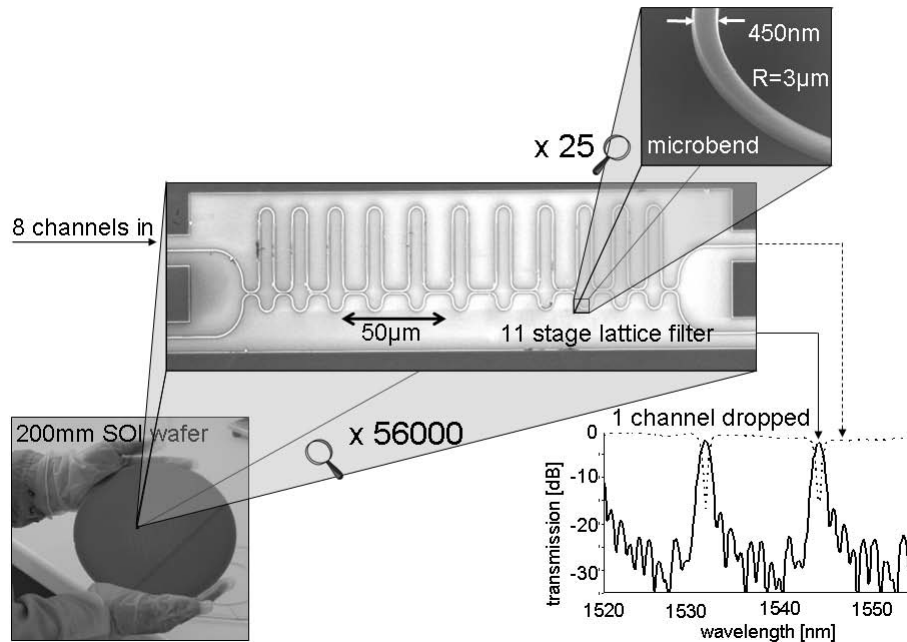
In this work, integrated optical filters were realized with wafer-scale technology, including deep UV optical lithography. The components are based on SOI photonic wire waveguides, confining light through total internal reflection. Several types of filters were implemented in order to get an overview of the possibilities and areas that need improvement. The filter types studied are well-known in lower index-contrast material systems: ring resonators, Mach-Zehnder based filters such as highly flexible lattice filters, and arrayed waveguide grating (AWG) wavelength routing devices.

The performance of the components was studied based on an in-house developed fabrication process. The process used throughout this work was optimized for general photonic structures of varying pitch and fill factor, and showed loss figures typically around 3.5 dB/cm. With an older process, not properly supporting periodic structures, wire waveguides showed a propagation loss down to 2.4 dB/cm.

We have demonstrated ring resonators with bend radii of just 1 to 5 µm. Single-ring channel drop filters with a large free spectral range, low insertion loss and down to 20 dB extinction in the throughput channel were demonstrated. Lower loss and lower distributed scattering would allow for higher-order filters to be designed and fabricated.

Mach-Zehdner lattice filters dropping one channel out of 4 or 8 were designed and fabricated with WDM channel spacings of 100 to 400 GHz. These filters can be designed with filter synthesis methods well-known from the digital filter world. The fabricated devices showed a low loss, bandwidths as designed and an extinction of the order of 10 dB (down to 15 dB in the dropped channel port). Analysis showed the crosstalk is due to both coupler variations and phase errors in the delay lines. These phase errors were already reduced by broadening the waveguides in the straight sections of the delay lines.

Wavelength routers based on arrayed waveguide gratings were designed and fabricated. Devices with 4 to 16 wavelength channels and 100 to 400 GHz channel spacings, showed an insertion loss of down to 2.5 dB and a channel isolation level down to 15 to 17 dB. To reach this, the refractive index contrast was lowered around the star couplers through a double-etch process, and the arrayed waveguides were made less sensitive to small dimensional variations by broadening them.

The uniformity between identical filters on a scale of a few tens of micrometre, a few millimetres and between different dies was investigated. The characterization results point in the direction of short-range stochastic deviations that are dominating any deterministic behaviour. The short-range uniformity is clearly better than the longer range uniformity though. Combined with the measured crosstalk levels and the process, wafer and mask characterization, this points to waveguide width variations on a length scale of some hundreds of nm to a few micrometer are the dominating contributions that need to be eliminated. However, there are indications that such variations are already partly existing in the lithographic mask.

Next to the improvement of component performance, major issues towards real-world application of integrated optical filters in SOI are the connection to fibre and the polarisation dependence. Coupling between submicrometre waveguides and fibre seems difficult at first. However, over the last years, coupling methods with high performance and good alignment tolerances have been demonstrated and are, probably counterintuitively, enabled by the very high index contrast. However, actually physically connecting the silicon chip to fibre in a stable, reproducible way

supporting many fibres while avoiding a too large size difference between the fibre connector and the component size is still an issue. In this work, an AWG based wavelength router in SOI was fibre pigtailed using vertical fibre coupling and a standard eight fibre array connector. The device was subsequently used in a technical demonstrator for a high-speed interconnect application. This optical interconnect functioned as a reconfigurable optical backplane interconnecting electronic boards at 10 Gbps. With the SOI wavelength router at its heart, characterization of the backplane showed a scalability of 4 to 9 ports. This clearly illustrates the limits but also the possibilities of SOI components fabricated with wafer-scale technology today. As this optical backplane was investigated for use in satellites in space, a preliminary radiation hardness testing of SOI components was performed. The result indicates that while the passive optical filters in SOI show a higher radiation-induced spectral shift than their low-index-contrast counterparts, the deviation will only become significant on missions of long duration or in stronger radiative environments, and can be compensated for by temperature tuning anyway.

The components were fabricated with wafer-scale processes that are developed for fabrication of a wide range of structures, and not tailored toward the specific structures studied. Despite the present non-uniformity between and within structures, a component performance decent enough for first application tests was demonstrated, with optical filters with WDM channel spacings fitting on a footprint of the order of just 0.01 to 0.1 mm^2.

High yield and reliable fabrication of packaged photonic integrated circuits based on wafer-scale technologies, with a complex functionality and a wild variation of the underlying physical structures is still some way off. But we have paved a little part of the way towards this goal by showing photonic wire based, ultra-compact integrated optical filters in SOI, using wafer-scale processing, delivering a nice performance for first application tests and a basis for further development.

English Text

Contents

Contents

Chapter 1

Introduction

This work is about guiding light through ultra-small wire-like structures of silicon. Based on these *photonic wires*, extremely compact wavelength filters for communication applications were designed and fabricated with electronic integrated circuit based technology.

Such a silicon photonic wire waveguide is shown in figures 1.1(a) and 1.1(b), guiding a light over a supporting oxide area. The largest fraction of the light is confined in an area of just 200 nm by 500 nm. This photonic wire can be bent with a radius of just a few micrometres without losing too much light, which enables extremely compact components to be designed. Functional devices with footprints smaller than $0.1\,\mathrm{mm}^2$ or even smaller than $0.01\,\mathrm{mm}^2$ can be built, as illustrated by figure 1.1(c). The key to this compact device size is the high refractive index contrast between the semiconductor core of the waveguide and its air and oxide cladding. This enables photonic integrated circuits (PIC) to become some orders of magnitude more compact than today's commercially available

Figure 1.1: A photonic wire cross-section, bend and component in silicon-on-insulator

components. While photonic integrated circuits were originally developed for use in the backbones of telecommunication networks, integration of a larger functionality on a smaller carrier is an enabling factor in moving photonics closer to the end user, in metro and access networks, for user applications such as sensing, or in size and weight-constrained environments.

Although silicon cannot yet be used to generate light with any significant efficiency, it is a good material for guiding infrared light with a wavelength longer than 1.11 µm, and available in high quality at relatively low cost. And, as important, it is the material of electronics, and more specifically CMOS components. Therefore, we can expect to use the huge technology base, maturity and high-yield processes of CMOS fabrication, and possibly to even combine photonics with electronics on a single chip. The structures shown in figure 1.1 are fabricated with CMOS based fabrication processes that were already developed at the start of this work [1, 2].

Basically, the structures are etched into a thin membrane of silicon, supported by a dielectric cladding, separating it from a silicon carrier wafer. This silicon membrane technology is based on silicon-on-insulator wafers, also used for CMOS. Over the last few years, experimental work has been performed on essential components for optical communication networks, such as modulators [3], detectors [4] and wavelength filters [5] in or based on silicon and silicon membrane technology.

It is known that the higher the refractive index contrast, the stronger the requirements imposed on accuracy. In most research environments, accurate but slow fabrication methods are used. CMOS based fabrication processes for silicon membranes have been used by a limited number of academic research groups only. A list is given in chapter 3. Real CMOS compatible photonics is just being shown commercially for some types of components and functionalities, but is still far off for general photonic structures.

A typical application for passive waveguide circuits are wavelength filters for use in telecommunication and datacommunication networks, where multiple laser wavelengths are used to transport different signals at the same time over the same physical carrier fibre. In this work, such wavelength filters are studied experimentally. Compared to lower refractive index-contrast materials, the device size is reduced by several orders of magnitude. Moreover, the devices are also excellent indicators of the quality of the fabrication technology.

The question whether silicon membrane based technology will eventually be able to deliver highly performing basic components (including

highly efficient sources) is a difficult one to answer, definitely if one poses the additional requirement of CMOS compatibility. This work tries to answer a small part of this question: *is it possible to obtain components for wavelength multiplexing and demultiplexing, fabricated with electronic IC technology, delivering good performance for telecommunication or emerging applications* ? The problem that arises is *how the current CMOS based fabrication processes perform, and what the fabrication tolerances are for proper functioning of wavelength filters*.

In the next section, high index contrast silicon-based photonic integration will be discussed, and the relation to electronics integration. Subsequently, the components studied in this work and the basic requirements of such components are introduced. These components are wavelength filters with telecommunication-oriented specifications. Finally, the objectives of this work are set out.

1.1   Silicon photonics integration

1.1.1   Integration

For components in optical networks, the cost of packaging and testing is as much as half and even up to 90% of the component cost. This is mainly due to the slow and labour-intensive steps needed to optically align fibres, integrated chips and discrete components with each other. By integrating more functionality on a single chip, the alignment and packaging cost can thus be considerably reduced. Additionally, integration through the very high index contrast components explored in this work may actually facilitate alignment, which could drive the cost down even more. A second reason for moving to higher integration levels is the need for smaller, lighter and highly functional devices, as photonics components move closer to the user, into the metro and access networks, and get new, consumer-oriented applications such as sensing which cannot do with bulky components.

1.1.2   High index contrast nanophotonics

A higher integration level can be obtained by using higher refractive index-contrast material systems. Physically, the advantage of a higher index contrast is the control over a wider range of propagation directions of the light compared to lower index-contrast materials. A direct consequence of this is the higher confinement of the light into the core material. The better control over light makes it possible that a change of propagation direction

of the light can be obtained on a very short length scale, with bends of just a few micrometre in radius. Also, components for splitting or combining light can be made much shorter. The higher confinement into the core is advantageous for active and non-linear processes, that can be used for generating, absorbing and amplifying light and for wavelength conversion, high-speed switching and other applications.

The higher the index contrast, the higher the integration level that can theoretically be reached. The lateral index contrast is important, but also a high vertical index contrast is, and allows for the optical power to be confined in an extremely small area, which is key to enhancement of non-linear and active effects. Therefore, semiconductor membranes suspended in low index contrast material seem a very attractive approach.

However, this comes at a cost: smaller dimensions are needed to keep waveguides single mode, and smaller gaps between waveguides are needed to allow weakly coupled structures with a limited length. High-index contrast structures are more sensitive to geometrical and refractive index deviations. Scattering losses scale with the index contrast through a power law. With feature sizes of tens to hundreds of nm, even nanometre-scale variations induce large variations in the resulting optical behaviour. This is discussed in chapter 2 and chapters 5-7. Therefore, the feature size should be extremely well controlled, down to the nanometre level. This is why one can speak about *nanophotonic* structures. In conclusion, we need processing technology with a better resolution, with smaller minimum feature sizes, and with a better control.

Although the sensitivity problem is fundamental, it is important to point out that a high index contrast, both horizontally and vertically, enables functionality that is simply impossible with a lower index contrast. This may even lead to the situation that some things are actually easier in a high index contrast system. A nice example of this is vertical fibre coupling through a high index contrast grating, which enables coupling to standard single mode fibre with a large alignment tolerance.

1.1.3   Silicon and CMOS compatibility

The drive for using technology from the electronics world for photonics is large. However, there is a difference between using CMOS technology or CMOS based technology and being CMOS compatible. This CMOS compatibility will be further commented on in chapter 3. Different levels can be considered: one can use fabrication tools from the CMOS world, CMOS based processes and CMOS compatible materials, one can use these CMOS fabrication tools in an environment where actual CMOS

fabrication is done, and finally one can use processes that are fully compatible with CMOS fabrication and circuits, mainly in terms of thermal budget during fabrication.

1.1.4   Parallel with electronic ICs

Often, especially for silicon photonics, a parallel is drawn between the history of integration improvement in the CMOS world and the photonics world. However, reaching the complexity and the reliability and yield of silicon electronics is a very difficult task for general highly integrated photonics circuits. This is mainly due to some important differences between electronic and photonic integrated components that cripple the analogy.

- **Never the same level of - a different kind of - integration** Photonic structures have a basic dimension (critical dimension or CD) that is somewhere in the order of the wavelength used. Also, in order to avoid crosstalk between waveguides, their minimum spacing will always be in the order of a micrometre. In the future, very interesting things will probably be done with subwavelength features, but still the global basic structures will be of the order of a considerable fraction of the wavelength, a few 100 nm. One notable exception is that of structures exploiting plasmon resonance and surface plasmon resonance (SPR) phenomena [6][1]. However, photonic and electronic integrated circuits have a different kind of integration. An arrayed waveguide grating such as those discussed in chapter 7 has a high level of functionality that would probably take the same chip footprint in its electronics counterpart.

- **No single material system** Silicon cannot do for photonics what it can do for electronics without the combination with other materials. While this is actively researched and will eventually be technically viable, it is a processing complication which makes every possible solution less elegant than the silicon-centric electronics world. (Although for several purposes, other materials are used or their use is being considered for electronics as well).

---

[1]A surface plasmon resonance is an electromagnetic wave bounded to the surface of a metal particle or layer. Plasmon excitations can have much smaller wavelength, and in this way extremely compact functional structures can be built, also for filtering applications at longer wavelength. However, the extremely high losses limit the number of such structures in an integrated circuit to just a few, and the SPR has a very limited bandwidth, which will again define the total chip real estate the global circuit needs

- **No single basic component**. In CMOS, components are built from NMOS and PMOS transistors, combined with some basic components as diodes, capacitors and resistors. In photonics, there is no such basic component as the transistor. Some researchers see the ring resonator as a candidate, but in the author's view, it is dubious whether this will become reality. Indeed, the ring resonator is a versatile component and many resonators together can be used to obtain the functionality of other component types, but those other components simply have so much advantages for some types of functionality that it will be very difficult for the economical decision of building everything with resonators to outweigh the consequences.

- **Analog versus digital**. In electronics, the functioning of the circuits as such is digital, based on the non-linear laws governing the underlying physical structures (transistors). While optical communication is digital, the functioning of many optical components is basically analog. An optical filter for instance acts similar to an analog electrical filter, although, as will be seen in chapter 4, for design purposes a very useful parallel between optical filters and digital electronic filters can be drawn. However, the analog nature of the implementation of the optical functions has a very important consequence on fabrication tolerances: while digital circuits can often live with relatively large variations, and more components can be added to obtain higher performance, this is much less the case for optical filters.

What kind of integration can be expected ? Two directions are possible. Either circuits evolve to a highly parallel functionality or a highly serialised functionality. Of course, these are extremes and circuits can end up somewhere in the middle. In parallel integration, a short chain of components, a short fibre-to-fibre light path for communication purposes, is duplicated many types on a single chip and a relatively low number of functions is applied to many ports and channels. This approach is being used for InP integrated components with commercial success [7], and is technically most feasible, as one can compensate for fabrication deviations between the parallel lightpaths by relatively simple tuning. In a more serial way of integration, a low number of channels or ports undergoes a large number of operations or some complex operations (complex or large components). In other words the input-to-output lightpath is longer. The main problem associated with this integration type is that every component in the lightpath needs to be aligned (spectrally, for instance), fabrication variations have a larger influence and more complex

tuning is needed. For integration with long lightpaths, insertion loss and crosstalk behaviour of basic structures and components such as filters are more important, as in general the same target levels of input to output loss and crosstalk will be specified. In electronics, there also is both strong parallel integration such as for memory, and more serial integration such as general processing units. However, even electronics chips with a highly serial functionality are built from physical layers with the highly parallel nature of arrays of transistor gates, metallisation and so on, with the fabrication process optimised for each layer. This is not (yet) possible for photonics fabrication.

Nevertheless, the differences between electronics and photonics integration do not degrade the relevance of using the huge silicon technology base for photonics.

1.2 Wavelength division multiplexing filters

What can integrated photonic circuits be used for ? Today, all telecommunication networks offering services like telephone, internet access and cable television are powered by an optical backbone. In the core of the regional network connecting larger cities, multiple wavelengths are used in order to maximise usage of the present optical fibres. Using this wavelength division multiplexing (WDM) technology, many thousands of voice or data connections can be carried over a single fibre[2].

1.2.1 Wavelength channel filters

In networks employing WDM, wavelength filters are essential components. These filters operate on wavelength channels, combining or splitting them from or into one or several outputs. At the optical communication channel endpoints, multiple channels are combined on a single fibre by a multiplexer or extracted from a single fibre by a demultiplexer. In a ring network, at each node a wavelength channel can extracted from the fibre by a *channel drop filter*, and data can be *added* on the same wavelength through the same or an accompanying filter. In a star network, at the interconnection between multiple ring networks, and in hybrid network topologies, wavelength routers act on multiple channels on multiple inputs and outputs simultaneously.

---

[2]More than 150000 phone connections can simultaneously be carried today by a single wavelength in an optical fibre

Such filters can be made in electronic form, with optical receivers and transmitters at the inputs and outputs. All-optical wavelength filters on the other hand do not need electro-optic conversion in order to perform their function, with the advantage of transparent operation. The main benefit is the possibility of insensitivity to modulation formats and higher-level communication protocol properties. These optical wavelength filters can operate passively, not consuming electrical power or only for temperature control purposes. With passive filters, static multiplexers, demultiplexers, channel drop and add filters, and wavelength routers can be obtained. Many other filters and structures can be built in passive waveguide circuits. For instance, wavelength filter types exist for equalisation of power levels between channels or for compensation of dispersion. Active structures can complement the passive functionality to provide wavelength conversion and regeneration of the signals, and could be integrated on the same chip, if the level of integration and the power and temperature constraints would allow.

Wavelength channel filter components exist today commercially, although the range of components is currently limited to a few types: fibre-based components, free space thin film filters, free space and integrated grating filters and integrated optical *arrayed waveguide grating* (AWG) multiplexers. The integrated optical components available only contain one or a few functional elements per chip and thus a complete system is bulky.

In the following, a review of filter specifications and filter types is given to situate this work. The reader not familiar with filter theory in general or optical filters will find a brief introduction in chapter 4.

1.2.2   Application areas

Wavelength division multiplexing is a technique developed for the backbones of telecommunication networks. Therefore, the components studied here are designed for operation around 1550 nm wavelength, have WDM compatible channel spacings and can have applications in telecommunication. However, other applications are possible.

In metro and access networks, optics are moved closer to the end user. In fibre-to-the-home (FTTH) schemes, often three wavelength carriers are used, for upstream and downstream communication and for broadcast. The bands used have coarse spacings in the 1550 nm band, or are divided over the 1300 and 1550 nm bands. However, WDM based technology could be used as well, and deliver more flexible networks and higher performance. An example of such an application is given in chapter 9.

Short-distance datacommunication mostly uses sources at shorter wavelengths, such as GaAs based VCSELs, often in combination with cheaper glass or plastic multimode fibre. However, the mature and advanced technologies of telecom could also be used for datacom applications: high performing, temperature stabilised, tunable lasers, and advanced detectors can be used to build board-to-board interconnect networks running at high speeds starting at 10 Gbps. Combining WDM technology can give massive channel bandwidths, or easy wavelength routing, or both, in such a network. Using SOI technology for such an application is a subject of chapter 9 as well.

Besides in optical communication networks, SOI wavelength filter technology can find applications in totally different fields such as sensing. By combining with MEMS or NEMS approaches, vibration and other sensors can be made. Due to their high temperature dependency, SOI filters are also well suited as temperature sensors. And through a change of the refractive index of the overcladding, chemical and biosensing is made possible. These applications are currently very actively researched.

1.2.3 Filter specifications

Wavelength filters for WDM applications have bandpass or bandstop characteristics and have specifications on the band centre frequencies, band width, band edge roll-off, isolation and band attenuation and ripple. The ideal filter for WDM applications has a large bandwidth within the allocated channel band, a steep roll-off, small in-band attenuation, and large out-of-band attenuation (small insertion loss, large channel isolation). For metropolitan networks, access networks and data communication applications, the requirements may be very different from those for long-haul telecommunication. Nevertheless, the mature laser and receiver components developed for telecommunication may also be used for other use, and the components studied in this work may find application there.

Insertion loss   The loss of the (de)multiplexed channels between entry and exit of the component (insertion loss) should be minimised, as only a limited amount of loss can be tolerated over a complete optical communication path. Loss arises from material losses, coupling losses between components and fibres, losses due to imperfect fabrication and losses intrinsic to the design of the component.

A second factor is the *loss non-uniformity*, as different channels may experience a different loss. Well-balanced power levels makes the life of

system designers easier and minimises the need for variable optical attenuators. Definitely in a highly integrated context with many components in a light path this is important as the loss non-uniformity of multiple components will add up in a complex way. Often the loss non-uniformity is intrinsic to the design of a component, like in an AWG wavelength router (Chapter 7).

Channel isolation A lack of channel isolation leads to crosstalk, which in its turn influences the optical signal power needed for communication with a low enough bit error rate (BER). While channel isolation is limited intrinsically by filter design, it turns out that for most photonic integrated circuits, the isolation is limited by fabrication variations. The larger the integration level, the more difficult it is to keep the isolation high.

The required isolation mainly depends on the power budget per channel, the aggregate losses and noise in the system, and the number of wavelength channels. The power budget is defined by the optical output power of the modulated source and the receiver sensitivity to arrive at a certain required BER. Normally, all noise is considered additive, driving down the signal to noise ratio (SNR). For a modest number of channels, crosstalk generally behaves additive. For a higher number of channels, a deviation from the additive behaviour can be noticed. The higher the number of channels, the higher the aggregate crosstalk in general, and the more strict the channel isolation requirement. In chapter 9, these considerations will be used to evaluate the scalability of the fabricated structures for one example application.

1.2.4   Filters studied in this work

In this work, three filter types are studied, the ring resonator, Mach-Zehnder based filters and the arrayed waveguide grating (AWG), representing a wide class of possible devices. The filters are designed with channel spacings of 100 to 400 GHz. This makes them compatible with other WDM technology, although the future application of the components will probably not be in the core networks. The work focuses on insertion loss and channel isolation, as these are influenced severely by fabrication imperfections and are to a large extent intrinsic to the material system. Bandwidth, roll-off and ripple are somewhat more depending on general design trade-offs relating to any technology, and although very important, it is of primary importance to limit the insertion loss and channel isolation first. In this work, mentioned insertion loss values can either indicate the loss from

input fibre to output fibre or the loss between input and output of the integrated optical component (waveguide to waveguide), which will be clear from the context.

1.3 Objectives

This work studies wavelength filters with WDM compatible channel spacings, based on high index contrast silicon-on insulator wire structures and fabricated with CMOS based wafer-scale technology. By looking at three filter types - the ring resonator, the lattice filter and the arrayed waveguide grating, covering three different filter design approaches, it is possible to obtain an overall view of what can be done using this technology. The work focuses on loss, channel isolation and reproducibility. It looks at the performance achieved with current state of technology, and what the future requirements are.

1.4 Overview

In chapter 2, a review of the properties of silicon-on-insulator photonic wires and passive wire circuits is given. Both the waveguiding properties of wires and their sensitivity to variations are considered. A number of basic circuit building blocks such as splitters and intersections are discussed, based on measurement results. In chapter 3, an overview of the fabrication using CMOS based processes is given. As the study of this and the process development was the work of Dr. W. Bogaerts [1], chapter 3 is not going into much detail. However, the possible fabrication variations are also discussed. Some general properties of wavelength filters and considerations on their sensitivity to fabrication variations are given in chapter 4. Also, some basics of general filter theory and the relation between optical and digital filters are discussed. In chapters 5, 6 and 7, ring-resonator-based filters, Mach-Zehnder based filters and arrayed waveguide gratings are discussed, respectively. The discussion is based on characterisation results, in combination with simulation in order to explain behaviour and predict or speculate on possible future performance. Then, based on the data in chapters 2 to 7, processing uniformity and accuracy are further discussed in chapter 8. Finally, in chapter 9 we show a practical application of nanophotonic circuits, a high speed, wavelength routed interconnect network, which demonstrates the current state of technology.

1.5   Original work, publications and acknowledgements

The original work performed by the author, described in this thesis, includes all simulation work, except where mentioned, design of ring resonators and ring resonator based structures, design of Mach-Zehnder, Mach-Zehnder lattice and Mach-Zehnder cascade structures, and various basic structures such as splitter trees, all static optical measurements mentioned in this work, except where mentioned, the interpretation of the static optical measurements, testing of the high-speed interconnect demonstrator and radiation hardness testing of SOI components (chapter 9).

Parameter and mask design of most AWG structures, including the polarisation-independent duplexer and the AWG router mask described in chapter 9, was performed by Dr. W. Bogaerts. The analysis of the wire loss measurements and the process development were also performed by Dr. Bogaerts [1]. SNOM measurements of ring resonator structures were performed by D. Néel at INSA Lyon.

Mask checking, correction and integration were done by J. Mees at IMEC. The actual processing was performed by the 200mm pilot line at IMEC.

An overview of all publications directly or indirectly resulting from this work is given in appendix C. Work performed but not directly described in this thesis led to co-authored publications.

The author helped opening up the fabrication facilities to other researchers in the group and European research groups, which led to a number of co-authored journal and conference publications.

In the work of Dr. G. Priem, and in the context of a collaboration with the NICT, Tokyo, silicon-on-insulator photonic wire and ring resonator structures were characterized for non-linear effects, leading to several co-authored publications with G. Priem, T.K. Liang and L.R. Nunes.

The author enhanced the functionality of the in-house developed CAMFR simulation tool to better support simulation of photonic crystal waveguide based structures, resulting in a co-authored publication with P. Sanchis (UPV, Valencia, ES)

Chapter 2

Photonic wires and circuits in Silicon-on-Insulator

In this chapter, the basic passive, linear properties of photonic wires and photonic wire circuits are studied, based on simple simulations and measurement results of basic components. First, the high index contrast wire in Silicon-on-insulator (SOI) is introduced. Before going into more detail, the simulation of wire structures and the characterisation methods are briefly discussed. Then, the propagation properties of SOI slab and wire waveguides are reviewed. Attention is paid to the sensitivity of structures to deviations from the desired structure, which will be used in chapter 4 and further chapters. Finally, measurement results of basic passive structures such as power splitters, MMI structures and waveguide intersections are discussed.

2.1 Introducing Silicon-on-insulator wires

The SOI material stack used throughout this work has a 220 nm thick crystalline Si core layer, separated from the Si substrate by a 1 µm thick amorphous SiO_2 bottom cladding (buried oxide, BOX). The top cladding is air or SiO_2. In some cases, structures were covered by polymer. By etching through the Si core layer completely, a very high horizontal index contrast is created. The high vertical index contrast thus not only introduces a high mode confinement in the vertical direction, but also facilitates achieving a high horizontal confinement. The waveguide geometry is illustrated in Fig. 2.1.

In the following text and all following chapters, an operation wavelength around 1550 nm is always considered. Most measurements were

(a) air top cladding   (b) oxide top cladding
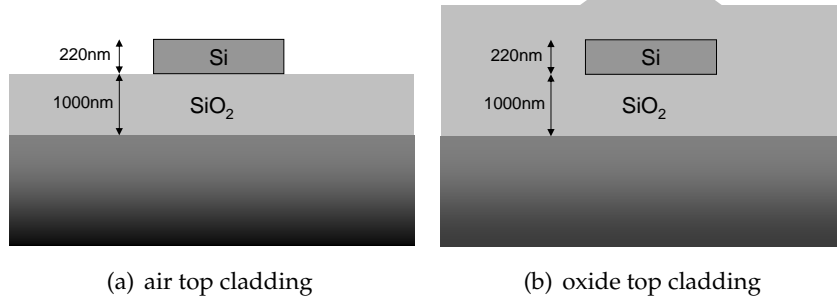
Figure 2.1: Drawing of the silicon-on-insulator wires used throughout this work.

	cSi	SiO$_2$
n	3.476	1.444
$\frac{dn}{d\lambda}$	$-7.6 \cdot 10^{-5}/nm$ [8]	$\approx -1.2 \cdot 10^{-5}/nm$
$\beta = \frac{1}{n}\frac{dn}{dT}$	$5.2 \cdot 10^{-5}/K$ [9]	$\approx 8.24 \cdot 10^{-6}/K$

Table 2.1: Physical properties of crystalline Si (cSi) and amorphous SiO2 at 1550nm wavelength and at 300K

performed within a 100 nm band from 1500 nm to 1600 nm. SOI allows for structures to work outside this band, but the dimensions and requirements on fabrication technology will of course be different. The material properties of silicon and silica around 1550nm are given in table 2.1. As a 220 nm Si slab thickness was used for the experimental work, all simulation results in this work are for this thickness, unless otherwise stated.

2.1.1   Why a high index contrast?

The advantage of a higher index contrast is the control over a wider range of propagation directions, larger propagation angles from the local propagation direction, or, reciprocally, a larger control over the wave vector space (k-space). A direct consequence of this is the higher confinement of the electromagnetic field in the waveguide core, as for similar longitudinal propagation constants β, the difference between β and the wave vector in the cladding is greater for a higher index contrast and thus the exponential tail of the mode is smaller. Both the greater control over k-space and the higher mode confinement have important advantages.

- A change of direction of light through bending or other refraction, through splitting or through diffraction can be realized on a shorter area. The bending radius can be as small as just a few micrometre. Splitters can be just a few micrometre long. Star couplers can be much smaller as a beam diffracts faster when freed from its horizontal confinement.

- Multimode structures with strongly coupled modes are more compact, as for the same waveguide cross-section there are more guided modes.

- Non-linear and active processes in the waveguide core are enhanced or have a larger effect. Therefore, structures exploiting non-linear or active processes are smaller.

2.1.2 Alternatives to photonic wires

While using strip waveguides to guide light is a relatively simple basic concept and can deliver powerful applications, other structures exist that can also be fabricated in semiconductor membranes and can give more design flexibility and possibly even more compact and higher performing components. For instance, photonic crystals, high index contrast periodic structures that guide and control light through distributed, omnidirectional reflection, are considered to deliver extremely compact devices with extremely flexible design options [10, 11, 1].

However, for passive, linear components, photonic crystal waveguides and structures offer little advantages over photonic wires. Often, their footprint is actually larger due to their distributed nature. Moreover, both simulation and fabrication are more challenging, and practically, losses are often higher due to fabrication variations. For non-linear and active applications however, photonic crystal structures do offer an advantage. Therefore, in the end both technologies are complementary. Actually, both structure types can be seen as extremes, the one isolated and the other one periodic with defects. More complex structures are being researched, with quasi-periodicities or very irregular topology, using optimisation algorithms [12, 13].

The fabrication technology used in this work was developed to support both isolated features such as photonic wires and periodic structures with varying pitches and fill factors. This has important consequences with respect to waveguide quality and structure reproducibility, as pointed out in chapter 3.

2.2   Simulation

Numerical simulation results in this and following chapters are obtained using the fully vectorial mode expansion software tools Fimmwave [14] and CAMFR [15]. Both have fully 3-D mode solvers based on variants of the film mode matching technique, but CAMFR uses a more analytic approach. While these tools can give great insight in the functioning of a structure, they have some limitations. Modelling of waveguide bends with very small radius and high index contrast is still problematic. Also, the calculation of scattering losses in three dimensional structures is still not straightforward. When simulating high-index contrast devices in full 3-D with mode expansion software, the structure is divided in slices with a constant cross-section along the propagation direction. However, a first problem is there is not always a clearly defined propagation direction and therefore a high number of modes is needed to reliably simulate the device. This increases computation time and memory requirements drastically, and decreases success rate. Secondly, there are often one or more slices in which not enough modes are found, failing the simulation of the whole structure.

Therefore, the mode expansion tools are only used to model the structure on the lowest level of waveguide cross-sections or simple propagation structures, or for 2D simulations using the effective index method. Most devices as such are modelled using their basic governing (semi-)analytical equations and circuit models, with input of material models and data from mode expansion software. For most wavelength filters, such a high level model can be written down, and some input of numerical methods such as mode profiles allows to obtain design guidelines and order of magnitude results. For the simulation of arrayed waveguide grating devices in chapter 7, a small software package was written, with material models and several approximate methods for simulating propagation through the device.

For a limited number of simulations, the FDTD tool Omnisim was used, which is much better suited for structures without a clear propagation direction. However, 3-D FDTD is very memory-hungry, limiting the computable domain in case of the structures studied here to about $10\,\mu m \times 10\,\mu m \times 2\,\mu m$ with a 30 nm grid. In the near future, a better resolution will be feasible through larger computer memories (>2GB) and computer clustering. Still, time-domain methods also have some inherent problems. For instance, simulating a very small ring resonator does not need a large amount of memory for the spatial dimension, but due to the high Q a very long simulation time is needed. Therefore, simulation of such a structure

(a) facet coupling   (b) vertical coupling
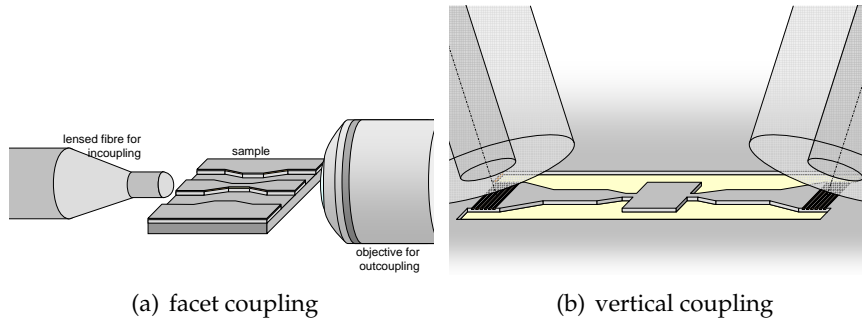
Figure 2.2: Coupling in and out of experimental structures using facet coupling and vertical coupling

on a single PC within acceptable time and with an acceptable spectral resolution was still impossible during the course of this work.

2.3 Characterisation

Most characterisation work in this chapter and the following chapters has been done through transmission measurements. As a light source, a tunable laser (external cavity laser) with a tuning range from 1500 nm to 1600 nm or a broadband light source is used. The amplified spontaneous emission (ASE) noise of an erbium doped fibre amplifier (EDFA) or a SLED (Superluminescent LED) deliver enough power around 1550 nm to characterise many structures in an efficient way. Light is launched from a polarisation controlled single mode fibre or polarisation maintaining fibre into the chip. At the output, light is collected and sent to a power meter in combination with the laser or spectrum analyser in combination with a broadband source.

2.3.1 End-fire method

One option is to couple laser light through cleaved facets in a broad access waveguide. A fibre with a lensed tip is used in order to couple efficiently as the core height is still only 220 nm. At the output, light is collected by an objective and imaged onto a detector. This is illustrated in Fig. 2.2(a). A linear taper converts between the broad access waveguides and the narrow wire structure to characterise. With uncoated facets, the facet reflection is of the order of 30%, but facet quality can vary. In particular, Si is hard to cleave and tends to cleave along multiple crystal planes. In order

to avoid feedback in the structures due to facet reflections, an AR coating can be applied. A HfO$_2$ coating was applied for characterisation of some ring resonator structures.

2.3.2   Vertical fibre coupling

For better alignment tolerances and more easy characterisation, most structures in this work were measured using vertical fibre coupling through grating couplers. These are one-dimensional gratings etched in a 10 μm broad waveguide. Detuned second-order gratings are used in order to avoid reflections at the waveguide facets. Light is coupled out to a fibre mounted under a small angle (8° to 10°) from the vertical. This is illustrated in Fig. 2.2(b). Linear tapers convert between the 10 μm broad waveguide with the grating coupler to a 3 μm broad waveguide and then down to the narrow wire structure. While the relaxed alignment tolerances are an advantage, small spectral variations over the wafer between the fibre couplers can make normalisation difficult.

A first generation of gratings was etched 40 to 50 nm deep, with a 610 nm pitch and uniform 50% fill factor. These have a coupling efficiency of about 20% and a 60 nm 3 dB bandwidth [16] without index matching material between grating and fibre. The second generation of couplers used has a 70 nm etch depth, a 630 nm pitch and a higher coupling efficiency of up to 35% when cladded with oxide, with an almost 60 nm wide 3 dB bandwidth.

2.4   Basic properties and light propagation

In this section, the basic propagation characteristics of light in our SOI structures is studied, based on simple simulations, literature and experimentally obtained data. The dispersion, confinement into the core, coupling between neighbouring waveguides, propagation losses and polarisation dependency are covered.

2.4.1   Silicon-on-insulator slab

With a 220 nm core height, slab waveguides in SOI have single TE and TM modes, as shown by Fig. 2.3. The single mode limit is 270 nm for asymmetric slabs and 250 nm for oxide-clad slabs. The TE$_0$ mode has an effective index of 2.83 for SOI with air top cladding and 2.85 with an oxide top cladding. Such thin SOI slab waveguides are very birefringent, with the

(a) air top cladding   (b) oxide top cladding
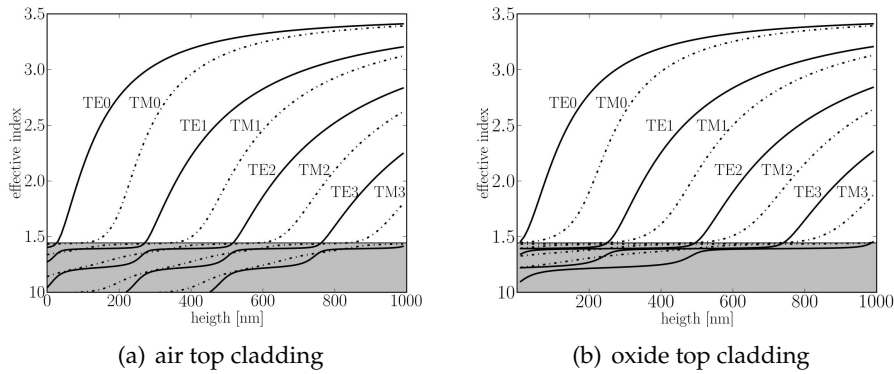
Figure 2.3: Simulation of the effective index of the slab modes as a function of core height. The gray area is above the light line.

TM mode having a much lower effective index. Only for a micrometre-sized Si layer, the slab becomes polarisation independent for the ground mode. A thicker silicon layer, between 300 nm and 400 nm, could be interesting for polarisation-independent waveguide devices (see Section 2.4.7), but the slab is then vertically multimode, and still polarisation dependent. Careful design is then necessary to avoid multimode problems in devices such as arrayed waveguide gratings [17].

For the TE_0 mode, the dependency of the effective index on small thickness variations is $3.5 \cdot 10^{-3}/nm$ around 220 nm thickness. The slab waveguide has a high dispersion, as for 220 nm thick slabs $\frac{dn_{eff,TE_0}}{d\lambda}$ is around $-5.5 \cdot 10^{-4}/nm$ for air and oxide top cladding, including material and waveguide dispersion. The group index therefore is around 3.7.

2.4.2   Photonic wire waveguides

The wire supports transverse electric (TE) like and transverse magnetic (TM) like modes. These modes are hybrid, with the TE modes having major E field components in the plane and the TM modes having their major H fields in plane. The level of hybridness depends on the actual cross-section, with the modes of square waveguides being most hybrid for symmetrically clad wires. In contrast to modes in lower index-contrast waveguides, the TE_{00} ground mode has a large E field component along the propagation direction.
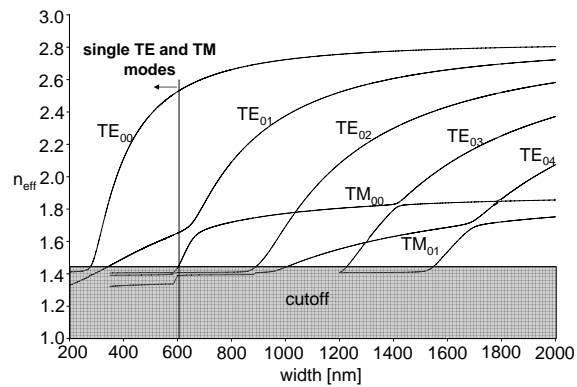
(a) air top cladding
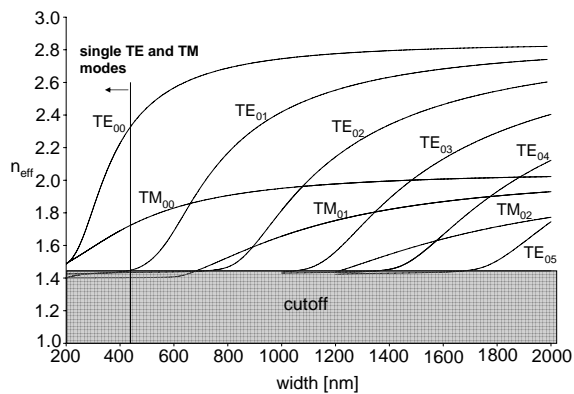
(b) oxide top cladding

Figure 2.4: Simulation of the effective index of the wire modes as a function of core width, at 1550 nm wavelength. The indices indicate horizontal mode numbers, as the waveguides are vertically monomode. Silicon layer thickness = 220 nm. They grey area is above the light line.

Figures 2.4(a) and 2.4(b) show dispersion diagrams for wires with air and oxide top cladding respectively, for a 220 nm thick Si core. At 1550 nm wavelength, the asymmetric wire has single TE-like and TM-like modes up to 600 nm width, and a single TE-like mode only below 350 nm width. However, with an oxide top cladding, the single-mode width is reduced to 430 nm, even though the index contrast is slightly reduced. For a 400 nm wide wire, the wire is single mode upwards of 1500 nm wavelength. In the asymmetric case, one can see that the TE-like and TM-like modes interact, making these modes completely hybrid at the anticrossing point.

2.4.3 Dispersion

The photonic wire has a normal[1] first-order dispersion in the width and wavelength region of interest. While the effective index is 2.44 at 1550 nm in a 500×220 nm^2 wire, the group index is 4.2. This is important when calculating the latency of an optical link, and must be used to calculate the delay lengths needed in filters to obtain a certain free spectral range.

The group velocity dispersion (GVD, second-order dispersion) of the TE$_{00}$ mode, including waveguide dispersion and a small contribution from material dispersion, is shown in figure 2.5(b). It was calculated at 1550 nm wavelength as a function of waveguide width. In the width range for single-mode operation, the dispersion is positive (anomalous) and reaches a maximum value of 580 ps nm^{-1} km^{-1}. This is 30 times higher than the dispersion in single mode fibres, but still only gives a propagation length limit of around 2 km at 10 GHz. Only at several THz, the maximum propagation length is of the order of cm and dispersion engineering does become important [18, 19]. Interesting features are the two zero-dispersion points, at 410 nm and 655 nm width, and the negative GVD for wider waveguides.

---

[1]There is often confusion about the sign and naming of first-order and second order dispersion. The n$^{\text{th}}$ order dispersion is *normal* when $\frac{d^{(n)}\beta}{d\omega^n} > 0$, with β and ω propagation constant and pulsation, respectively. Dispersion is called anomalous if $\frac{d^{(n)}\beta}{d\omega^n} < 0$. For the first order dispersion, normal dispersion therefore means $\frac{dn_{eff}}{d\lambda} < 0$, or $n_g > n_{eff}$: the phase velocity is larger than the group velocity. Anomalous first order dispersion has $n_g < n_{eff}$. Normal second order dispersion (frequently called just dispersion) has $\frac{d^2 n_{eff}}{d\lambda} > 0$, or $\frac{dv_g}{d\lambda} < 0$: a negative group velocity dispersion (GVD). Anomalous second-order dispersion = positive GVD.
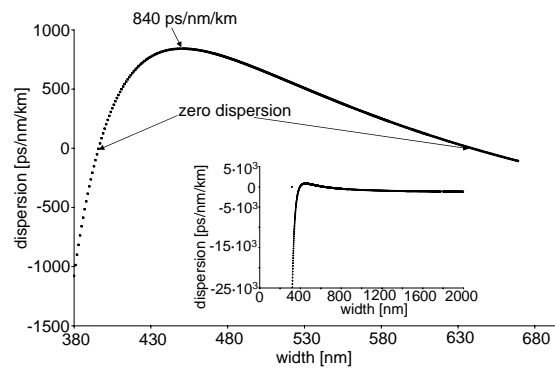
(a) air top cladding
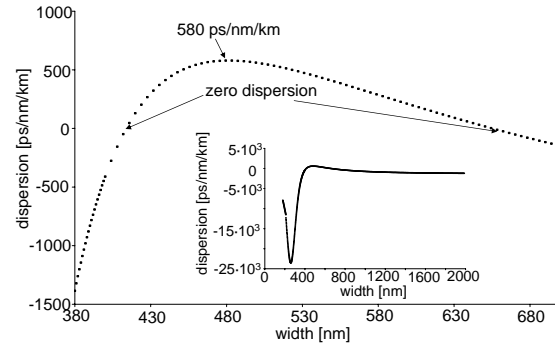
(b) oxide top cladding

Figure 2.5: Simulated dispersion of an oxide cladded SOI wire as a function of waveguide width, at 1550 nm wavelength. Simulation with material dispersion included.

2.4.4 Confinement

Two frequently used measures for the confinement are the mode effective core area, which basically is the integrated intensity, normalised over the intensity profile[2] and is used to model non-linear and other effects [20], and the filling factor, which is the fraction of power in the core[3]. The effective core area and the filling factor of wires with air and silica top cladding are given as a function of waveguide width in figure 2.6. The wire with air top cladding has a minimum mode effective core area of $0.17\,\mu m^2$ around 450 nm width. The filling factor is 74% in that case. With an oxide top cladding, the minimum effective core area is just slightly larger, at $0.19\,\mu m^2$ around 550 nm width, with a filling factor of 76%. The filling factor rises to 80% above 1000 nm. The vertical mode size, calculated from the electric field intensity at the vertical mirror plane through the wire, is given in figure 2.7. Clearly, for widths from 400 nm onwards, the mode height is not changing much. There is a small vertical mismatch between narrow and wide waveguides, which can give rise to loss and reflections. Simulation in Fimmwave results in a -30 dB reflection for a 400 nm wide waveguide incident on a slab.

2.4.5 Coupling

The integration density is to some extent defined by the minimum waveguide centre-to-centre pitch necessary for avoiding crosstalk between waveguides. For instance, in arrayed waveguide grating devices (chapter 7) with a large free spectral range, the size of the waveguide array is defined by the minimum pitch. Simulation of the coupling should be done in 3-D, as both the mode tails and the vertical dimension are important. Figure 2.8 shows the resulting minimum pitch for wires with oxide and air top cladding. The minimum pitch for a coupling less than -20 dB over a cm length was calculated. The pitch goes through a minimum, as the confinement saturates for wider waveguides. As can be expected, the minimum pitch is higher for oxide top cladding and for longer wavelengths. With an oxide clad waveguide at 1600 nm wavelength as a worst case, a gap of 1 μm is sufficient beyond 700 nm wire width. For more narrow wires, a 1.5 μm gap is a safe choice.

---

[2]The mode effective core area is defined as $\left(\int I\left(x,y\right)dxdy\right)^2 / \int I\left(x,y\right)^2 dxdy.$

[3]The filling factor is defined as $\int_R P_z\left(s\right)\cdot ds / \int_\infty P_z\left(s\right)\cdot ds.$

(a) air top cladding

(b) oxide top cladding

Figure 2.6: Simulated effective mode area and filling factor for wires at λ=1.55 μm

Figure 2.7: Simulated mode height of an oxide-clad wire as a function of width, for λ=1.55 µm, along the mirror vertical plane of the wire.
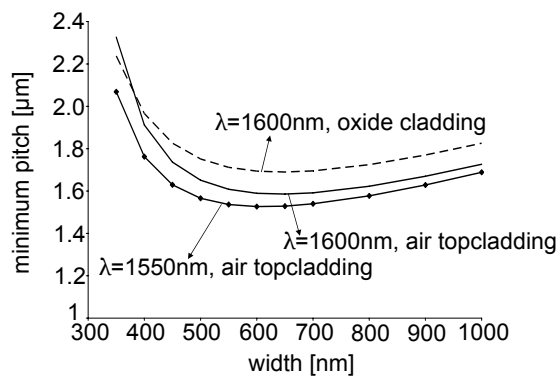
Figure 2.8: Simulated minimum pitch needed for less than -20dB/cm coupling between parallel wires.

2.4.6   Propagation losses

Propagation losses in straight SOI wires arise from multiple sources:

- **Bulk absorption** Pure crystalline silicon has a negligibly low absorption around 1550 nm. The Si top layer of the used wafers is however lightly Boron-doped, introducing a small absorption. From the resistivity specs of the Si layer, 8.5 to 11.5 $\Omega \cdot cm$, and based on the resistivity of bulk silicon as a function of Boron density [21], a doping level of $1.5 \cdot 10^{15}$ can be deduced. Extrapolating the experimentally obtained absorption as a function of free hole density [22], and given the fact that for low doping levels the free hole density equals the Boron density, the (linear) absorption is estimated to be around $4 \cdot 10^{-3}/cm$ or $0.02\,dB/cm$, not taking sheet effects on the resistance into account.

- **Surface state absorption** Dangling bonds and bonded elements at the sidewall introduce absorption centres. Generally, there are dangling bonds at the interface between the silicon and the thermal oxide top atom layers. Borselli et al. showed this source of loss can be significant [23], but is not easy to quantify experimentally. As this study was for microdisks, it is not clear what the actual impact on photonic wires would be.

- **Nonlinear absorption** If optical power is high enough, non-linear effects start to become important. Especially in highly confined photonic wires and wire-based resonators, this can already happen at relatively low peak power levels in the mW to W range [24]. The two major absorption effects, each accompanied by refractive index changes and thermal effects, are two-photon absorption (TPA) and free carrier absorption (FCA). At high optical powers, two photons can be absorbed together by an electron that then gets into the valence band. The free carriers that are generated in this way, can in their turn absorb photons and become hotter. As the free carrier absorption is proportional to the free carrier density, and the free carrier density here results from linear absorption or TPA, this is a non-linear FCA process. In this work, the power levels in the input wires are low enough, even in resonators, so that non-linear absorption will not be noticed. However, the power levels are high enough for secondary temperature effects to be measurable.

- **Substrate leakage** As the refractive index of the waveguide core is the same as that of the substrate, only leaky and not truly guided
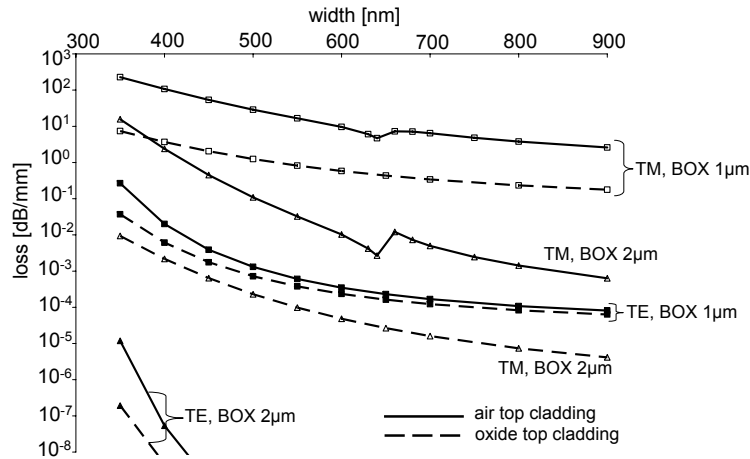
Figure 2.9: Calculated substrate leakage as a function of wire width.

modes will exist in the wire. During propagation, part of the power is coupled to radiative modes in the substrate and lost. As the tail of the mode into the cladding is exponentially decreasing with depth in the cladding, this substrate leakage can be expected to decrease exponentially with cladding thickness. Therefore, substrate leakage can be avoided by using a sufficiently thick BOX. The substrate leakage was calculated with CAMFR, using transparent boundary conditions. Figure 2.9 plots the calculated substrate leakage for the TE_{00} and TM_{00} mode as a function of wire width, for a BOX of 1 µm and 2 µm and for air and oxide top claddings. For the TE_{00} mode, with a 1 µm BOX the substrate leakage only attains measurable values for wire widths below 400 nm. The TM_{00} mode however has a much larger loss due to its lower vertical confinement. With a 2 µm BOX, the losses are negligible for all widths for TE_{00}, and are greatly reduced for TM_{00}. The oxide clad wires always have a lower substrate leakage, as the mode profile is pulled more into the upper cladding. W. Bogaerts calculated the substrate leakage for 220 nm thick SOI slab waveguides, arriving at 10^{-4} dB/mm for TE_0 with a 1 µm buffer [1, 25].

- **Rayleigh scattering** Small, deep sub-wavelength index perturbations in the bulk of the waveguide give rise to scattering following a λ^{-4} law in bulk and low index contrast waveguides. Rayleigh scattering is a fundamental limit in dielectric waveguides such as fibres. In

high index contrast crystalline semiconductor waveguides, the number of scattering points can be expected to be much lower, and a significant part of the scattered light can be expected to couple back to the guided mode. Because of the high quality SOI material used, the Rayleigh scattering can be expected to be extremely low. Nevertheless, it is the most basic lower limit of propagation losses.

- **Scattering at sidewall roughness** Light scatters at irregularities both on the sidewalls and top and bottom walls of the core, which is a major cause of propagation loss in photonic wires. Roughness on straight wires gives rise to scattering for spatial frequencies of the roughness between $\beta - n_{clad}k_0 < \xi < \beta + n_{clad}k_0$ [26], or $3.9/\mu m < \xi < 15.6/\mu m$ in the SOI wires studied here. In other words, periodicities in the roughness on length scales larger than 260 nm do not contribute to scattering loss. However, as we will see, these spectral components have a large influence on phase deviations of the propagating mode. Theoretical calculations of propagation losses due to scattering at sidewall roughness have been performed in the fully three-dimensional (3-D) case only very recently [26]. The authors show that previous two-dimensional (2-D) calculations can seriously overestimate propagation losses, as the losses depend very much on the waveguide cross-section. The dependency of scattering losses on wire width is not straightforward. Interestingly, the calculations show that for Si wires with a 220 nm core height, embedded in silica, the TM-like mode suffers a scattering loss which is an order of magnitude smaller than that for the TE-like mode.

  The losses depend on the roughness statistics. Often, the roughness is described by an exponential correlation function with a given correlation length. In [26], scattering loss for silicon wires with a silica cladding were calculated in the single mode region for a correlation length of 50 nm. Normalised on the variance of the roughness amplitude, the result for the wires studied here is 0.6 to $0.7\,dB/cm/nm^2$ for the TE$_{00}$ mode and about $0.1\,dB/cm/nm^2$ for the TM$_{00}$ mode. For other correlation lengths, a correction factor needs to be applied. Table 2.2 shows the scattering losses for the wires studied here, derived from [26], taking an approximate correction factor for the correlation length into account.

- **Scattering at top surface roughness** The top silicon surface also has some roughness, which is mainly due to the polishing process used during fabrication of the wafer. The scattering arising from this rough-

$\sigma_{roughness}$	L_c	TE-like	TM-like
2 nm	20 nm	1 dB/cm	0.16 dB/cm
3 nm	20 nm	2.2 dB/cm	0.36 dB/cm
2 nm	50 nm	2.4 dB/cm	0.4 dB/cm
3 nm	50 nm	5.4 dB/cm	0.9 dB/cm
2 nm	150 nm	3.6 dB/cm	0.8 dB/cm
3 nm	150 nm	8.1 dB/cm	1.8 dB/cm

Table 2.2: Approximate simulated losses due to scattering at sidewall roughness, for SOI wires with 220 nm core height and 400 to 470 nm width, for different roughness standard deviations and correlation lengths. Based on simulations results in [26]

ness gives a loss limit for waveguides, which is defined by the wafer fabrication and thus beyond our control, unless the wafers or waveguides are treated, for instance through chemical-mechanical polishing (CMP) techniques before further processing. In chapter 3 an AFM measurement on an SOI wafer is shown. Within a 2 µm × 2 µm measurement window, the rms roughness is about 0.1 nm. To calculate the scattering, the statistical behaviour of the roughness is important. In the AFM image (Figure 3.10), the roughness appears in typical dots and holes of some tens of nm in size. Correlation analysis in chapter 3 shows a reasonably well fitting exponential correlation function with a correlation length of 100 nm. Calculating the scattering loss of an SOI slab waveguide through the formula of Payne and Lacey [27], we arrive at a maximum loss of 0.043 dB/cm, unfortunately exactly at a correlation length of 100 nm. Nevertheless, this loss figure is very small.

Experimental results

We have measured propagation losses of wires by comparing straight wires of different lengths. The contrast of the Fabry-Perot fringe spectrum resulting from the reflection of cleaved facets is used to accurately calculate the propagation loss [1]. Table 2.3 lists the results. The losses are strongly dependent on waveguide width, as can be expected from the difference in confinement and field strength on the sidewalls for different widths. For 500 nm width, a propagation loss of only 2.4 dB/cm is achieved, which is already low enough for many experiments. This value was obtained with an early fabrication process, as described in chapter 3. A second process

Width	Propagation loss
400 nm	3.4 ± 0.17 dB/mm
440 nm	0.95 ± 0.18 dB/mm
450 nm	0.74 ± 0.09 dB/mm
500 nm	0.24 ± 0.16 dB/mm

Table 2.3: Measured propagation loss of the TE-like mode in straight wires as a function of waveguide width. Reproduced from [1]

(see chapter 3), which was further used throughout this work, yields a higher loss, with measured values consistently around 3 to 3.5 dB/cm.

For the TE-like mode, we see an exponential dependency on waveguide width. The magnitude of the difference cannot be explained by substrate leakage, and the simulations in [26] indicate this dependency cannot be explained by scattering at sidewall roughness if the roughness is equal for different wire widths. In order to contribute the measured losses for the TE-like mode at 450 nm width fully to scattering loss, $\sigma_{roughness}$ should be of the order of 4 to 5 nm at correlation lengths below 50 nm, or about 3 nm for longer correlation lengths. Current sidewall characterisation results are inconclusive on this. All this points in the direction of a larger absolute roughness amplitude (or a different statistic) for narrower wires, or another loss mechanism, which could be found in sidewall absorption, or a combination of both.

As the confinement of the TM-like mode is much lower than that of the TE-like mode, substrate leakage loss is higher, and a higher propagation loss results. With a thick enough buffer layer however, the simulations by Barwicz et al. [26] indicate the TM polarisation would suffer less scattering losses.

Propagation losses have been reduced over the past years to a few dB/cm by several groups. An overview is given in Table 2.4. These losses are still much higher than in silica waveguides or even large SOI rib waveguides, but as the minimum bend radius is reduced by orders of magnitude too, the higher propagation loss does not give a higher loss per component. In complex integrated circuits with many components and long total interconnection lengths, this picture changes. Therefore, losses need to be reduced further.

Affiliation	cross-section [nm^2]	Loss [dB/cm]
IMEC[2]	500×220	2.4
IBM[28]	445×220	3.6
Cornell[29]	470×270	5.0
NTT[5, 30]	400×200	2.4
Yokohama[31, 32]	400×320	105.0
MIT[33]	630×220	3.4
LETI[34]	300×300	15.0
	500×200	5.0
Columbia[35]	600×260	110.0
NEC[36]	300×300	18.9
Aachen/AMO[37]	500×450	1.9 (in bend)

Table 2.4: Measured waveguide losses reported by several groups

2.4.7 Polarisation dependency

As evident from figure 2.4, photonic wires are extremely polarisation dependent. It is possible to obtain equal propagation constants for the TE$_{00}$ and TM$_{00}$ modes by using a square core waveguide. The single mode dimension limit for a square core is $w = h = 320\,nm$. The linewidth needs to be reduced and thus the resolution of the fabrication process must be high enough. The main problem is however that the modes are very sensitive to small dimensional inaccuracies, and in a different way. It is possible to use stress-induced anisotropy in order to compensate for small birefringence in large core SOI waveguides [38], but this is much less a solution in the case of detrimental behaviour due to nm-scale variations.

Therefore, it is extremely difficult to obtain real polarisation independence within a single SOI wire circuit. An alternative is polarisation diversity by splitting up the polarisations at the input, sending them through two separate circuits, and combining them again at the output. With a two-dimensional extension of the vertical fibre coupler, this can be done in a very elegant way [39]. The grating couples the orthogonal polarisations to separate waveguides, but from the point of view of these waveguides, both polarisations are TE, and two identical circuits can be used. At the output, a second grating coupler couples both polarisations back to the fibre. This approach has two major problems: the circuit size is doubled, and a very good fabrication accuracy is needed so that the two circuits are exactly identical. Still, with two separate structures, transmission can be optimised for both polarisations by simple tuning, whereas with a square-

(a) height variation   (b) width variation
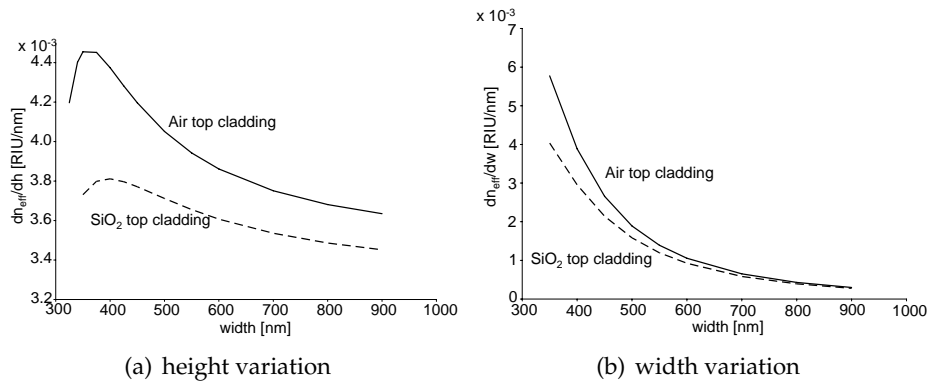
Figure 2.10: Simulated sensitivity of the effective index of TE_0 to waveguide height and width.$\lambda = 1.55\,\mu m$.

core approach a complex birefringence control is needed in every part of the single circuit.

2.5 Sensitivity

It is well known high index contrast waveguides are very sensitive to small variations in dimension or refractive index. Here, we quantify this sensitivity for the TE_{00} mode of wires with a 220 nm Si height. Also, the influence of random sidewall variations is considered. Due to the much lower confinement of the TM_{00} mode, its sensitivity to variations will be much lower. However, structures cannot be made as compact as with the TE mode. Of course, this changes with a changing waveguide form factor.

Figures 2.10(a),2.10(b),2.11(a) and 2.11(b) give the change of effective index of the TE_{00} mode for small variations of the waveguide height, width, core refractive index and top cladding refractive index, respectively. It is clear the waveguides are extremely sensitive to nm-scale dimensional variations. This sensitivity improves for wider waveguides. A change of refractive index also considerably influences confinement. Therefore, the effective index change is 15% to 20% larger than the refractive index change for wires with air top cladding , and maximum 5% larger for wires with oxide cladding, for single-mode wire widths. As confinement in the core is high, the sensitivity to the cladding index is smaller than to the core refractive index. For narrower waveguides, confinement gets smaller and the sensitivity higher.

(a) core index

(b) cladding index

Figure 2.11: Simulated sensitivity of the effective index of TE_0 to core and top cladding refractive index. $\lambda = 1.55\,\mu m$. In (b), the influence for different cladding indices are shown.

Temperature Si photonic wires are also sensitive to temperature. Combining figure 2.11(a) with the data in table 2.1 gives an effective index dependency on the temperature of about $1.9 \cdot 10^{-4}/K$.

Stress Strain induces changes in the electron density and in the average band gap in Si [40], causing a change in refractive index. The influence of stress changes is difficult to quantify. During fabrication of the SOI wafer, the wafer bonding results in a stressless material stack at the bonding temperature of 1100°C. Cooling down the wafer introduces in a compressive stress in the oxide and tensile stress in the silicon. However, the structure is partly relaxed by wafer bow. By etching the structure, the stress in the core layer is further influenced. Silicon has a {100} Young's modulus of about 130 GPa and a photoelastic tensor with a main p_{11} element with a magnitude of the order of 0.1 [40]. A stress change will give rise to a refractive index change in bulk Si of the order of $\frac{dn}{dP} = 1.4 \cdot 10^{-8}/kPa$. A stress change induced wavelength shift will therefore only rise to pm levels for stress changes of the order of some 100 kPa. For large cross-section SOI ridge waveguides, it is possible to control the birefringence of the waveguide by inducing stress in the core with a silica overcladding. This can induce a stress of the order of a few 100 MPa [38]. In photonic wire filter components, such stresses would induce wavelength shifts at the nm level.

2.5.1   A stochastic model for sidewall variations

Sidewall width variations (and wire height variations) are stochastic in nature, with some degree of correlation between variations. As sidewall roughness frequently appears as a curtain shape, it is often described as a one-dimensional distribution d(z) with a zero mean (the mean edge position). The distribution is considered Gaussian, and the stochastic process is characterised by its correlation function

$$R(l) = < d(z)d(z+l) > \tag{2.1}$$

where $<>$ indicates the ensemble average over the possible outcomes for a certain length of wire. Calculations of scattering at roughness [26] or phase errors are then based on the correlation function. This correlation function is often considered exponential or Gaussian, based on a limited number of measurements however. For a more complete description, a more accurate model on length scales shorter than the correlation length is to take the fractal dimension into account, describing the line or surface using fractal geometry [41]. Also, the correlation function is often more complex, with peaks around several characteristic distance differences (corresponding to peaks in the spatial frequency distribution) [42]. Chapter 3 goes into more detail on the sources of sidewall variations.

For the simulations in chapters 6 and 7, a simple model is used that allows to look at the influence of one-dimensional roughness with a given σ and correlation length L_c of a Gaussian or exponential correlation function. This model was proposed and used by Goh et al. to model phase errors in silica-based waveguides [43]. The model has its limitations when several typical spatial frequencies corresponding to quasi-periodicities are present in the roughness, but can still be used to determine the order of magnitude of each of these.

Within the correlation length, the sidewall position does not change much. Therefore, the rough waveguide is modelled as a series of constant-width stubs with a length L_c. This is shown in figure 2.12. The system is considered single mode and lossless, without coupling between TE and TM modes. It is therefore only valid for long enough L_c, as roughness-type variations give rise to scattering. Let $\sigma_{d(sidewall)}$ be the standard deviation of the waveguide edge with respect to the mean edge position. For uncorrelated edges with an equal distribution, the variance of the waveguide width is given by $\sigma_w^2 = 2 \cdot \sigma_{d(sidewall)}^2$. With this model, the variance of the phase deviation through propagation of a piece of waveguide with length L is given by (see appendix A)

$$\sigma^2 \left(\delta\phi \right) = \begin{cases} 2 \cdot \left(\frac{\partial \beta}{\partial w} \right)^2 \cdot \sigma^2_{d(sidewall)} \cdot L^2 & , \quad L < L_c \\[3mm] 2 \cdot \left(\frac{\partial \beta}{\partial w} \right)^2 \cdot \sigma^2_{d(sidewall)} \cdot L_c \cdot L & , \quad L > L_c \end{cases} \qquad (2.2)$$

For $L > L_c$, the standard deviation of the phase deviation is therefore proportional to $\sqrt{L}$. Figure 2.13 gives the phase error as a fraction of $\pi/\sqrt{\mu m}$ as a function of the correlation length and σ_d, for different widths. For 400 nm width, a phase error of $\frac{\pi}{100}\mu m^{-1/2}$ is reached for correlation lengths longer than 10 µm for a roughness with σ_d about 1 nm. For wider waveguides, a σ_d of 5 nm is needed to reach the same phase error for L_c=10 µm. Even so, it is clear 1 nm scale width variations on a longer length scale along the propagation direction are extremely detrimental. The influence of roughness with correlation lengths that generate scattering loss is smaller, but still a σ_d of only 5 nm can cause phase errors in the $\frac{\pi}{300}\mu m^{-1/2}$ to $\frac{\pi}{100}\mu m^{-1/2}$ range for smaller waveguides. However, for wider waveguides the influence of short length-scale roughness is greatly reduced.

Figure 2.14 gives a better view of the wire width dependency for some correlation lengths. For a correlation length of up to a few tens of nm, typical for the curtain-shaped roughness often appearing (see chapter 3), widening waveguides from 450 nm to 800 nm reduces the phase error width half an order of magnitude. For a correlation length of 500 nm, which could represent *wavy* waveguide sidewalls, the effect of widening the waveguide is smaller, but still, it means a reduction of a catastrophic $\frac{\pi}{50}\mu m^{-1/2}$ to about $\frac{\pi}{200}\mu m^{-1/2}$, which would be acceptable for short waveguides.

Figure 2.12: Phase error model for stochastic width variations. The waveguide is considered to have a constant within chunks of a correlation length L_c. Both sidewalls are uncorrelated.

(a) width=400 nm

(b) width=500 nm

(c) width=800 nm
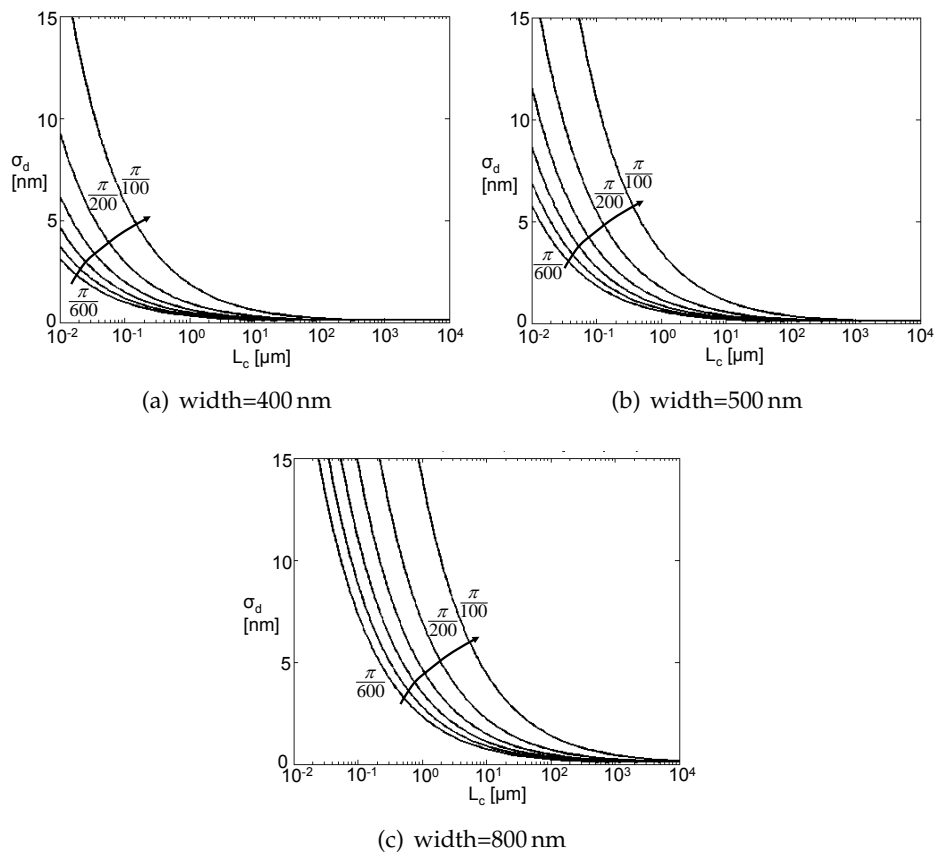
Figure 2.13: Simulated standard deviation of phase error per root micrometre ($\sigma_{\delta\phi}/\sqrt{\mu m}$) as a function of correlation length and standard deviation of roughness amplitude, at λ=1.55 μm
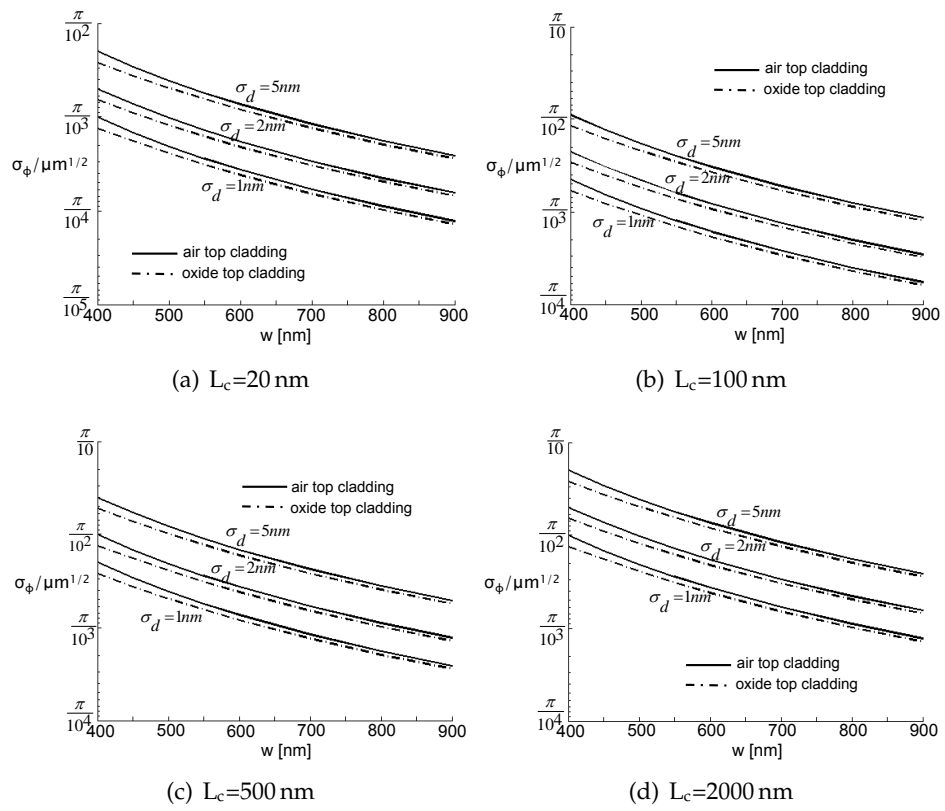
Figure 2.14: Simulated standard deviation of phase error as a function of wire width, at λ=1.55 µm $(/\sqrt{\mu m})$

In chapters 6, 7 and 8 these data will be used to interpret measurements and predict possible performance.

2.5.2   Conclusion: nanophotonics

While the waveguides have dimensions of several hundred nm, the phase velocity of their ground mode is highly dependent on the waveguide dimensions. With variations on a scale of 1 to 10nm, the effective index change is as large as 10^{-3} to 10^{-2}. Therefore, the accuracy on the lateral dimensions should be better than 1 nm in order to avoid phase deviations during propagation. As we will see in section 4.3.2, such phase deviations have a large influence on the spectral response of wavelength filters. Because of this need for accuracy, we can speak about nanophotonic waveguides and structures.

2.6   Basic structures

As the main reason for using nanophotonic SOI structures is the higher integration level, we need to have an idea about the performance of simple structures such as bends, waveguide intersections and power splitters. The trade-offs governing their performance will determine circuit footprint, together with those of the actual functional components such as wavelength filters. In the following, some measurement results of such structures are discussed.

2.6.1   Bends

Apart from the intrinsic bending loss, both the substrate leakage and the scattering loss are influenced by waveguide bending, as the mode is slightly pushed outwards. However, it is not directly qualitatively clear what this influence is. Bends in our current SOI wires also suffer polarisation conversion loss.

Bend losses were measured using spiral waveguides with various bend radii. Spirals were cascaded in a single structure in order to arrive both at significant wire lengths and a large number of bends. For each radius, eight spiral structures were measured covering the (length,number of bends) parameter space in an efficient way, with a total length up to 50 mm and up to 550 bends. Figure 2.15 shows measured loss through the spiral structure as a function of total length and number of bends, for a 560 nm wide waveguide. The strong dependency on the bend radius is

Figure 2.15: Measured loss through a spiral as a function of total length and number of bends, for different radiuses. Wire width is 560 nm, $\lambda = 1.54$

clearly visible. Both the straight wire loss and the excess loss of a 90° bend can be extracted from this measurement from the slopes along the axes of a bilinear fit through the data. Figure 2.16 shows the resulting excess loss as a function of bend radius.

A simulation of bend losses was performed by L. Prkna (formerly at Prague University) with a home-developed 3-D bend mode solver (BMS-3D) based on the film mode matching method [44]. Horizontally, the structure is left unbounded, while vertically, the structure is enclosed between electric or magnetic walls. Another simulation was done with Fimmwave using a very similar method with transparent boundaries at the left and right side, and reflecting walls at the bottom and top. Fimmwave under-

Figure 2.16: Measured excess bend loss as a function of radius, for the spirals from figure 2.15

estimates the losses for larger radius, and fails for smaller radii. BMS-3D predicts a radius with a minimum of loss per 90°, as can be expected. Despite the lack of accuracy for smaller radius, the intrinsic bend losses predicted are very small.

Simulation of the interface between a straight and bent waveguide in Fimmwave shows the interface loss is dominant for a large range of radii, and the main source of loss is coupling between the TE-like and TM-like modes at the interface. The simulation result is shown in figure 2.17(b) for different waveguide widths. As the TM-like mode experiences a (much) larger loss than the TE-like mode, the net effect is an interface loss. With a thicker buried oxide, the influence of TE/TM coupling may be seen in devices. Simulation in Fimmwave shows that transmission can be marginally improved by a very small offset of 5 nm between the straight and bent waveguides, as is used for lower index contrast waveguides. The transmission to TM is then also reduced. However, the improvement is not drastic. The simulations were done with straight sidewalls. Sloped sidewalls may give a higher loss and polarisation conversion. Moreover, with sloped sidewalls, the polarisation conversion through a straight-bend-straight interface will depend on the angle and radius of the bend. Polarisation conversion in silicon wires was studied in detail by Sakai et al. for smaller radii [45].
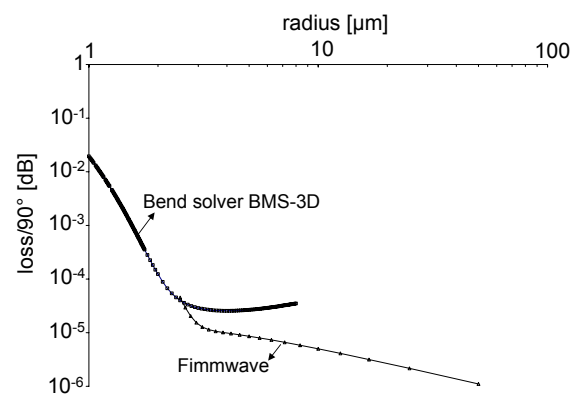
(a) Bending loss

(b) Straight-bend interface transmission and reflection. Modes with tilde indicate bend modes.

Figure 2.17: Simulated radiation and interface losses in bends. Radiation losses simulated using the BMS-3D solver [44] and Fimmwave with transparent left and right boundaries. Interface losses simulated with Fimmwave.
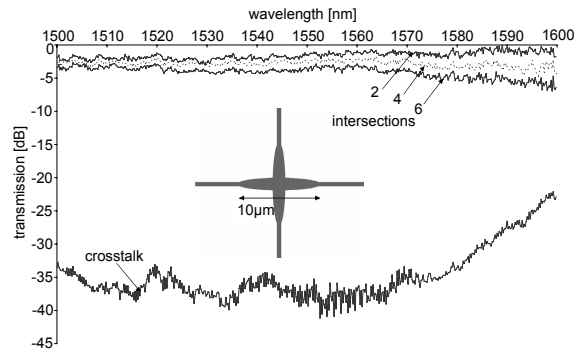
Figure 2.18: Normalised measured transmission and crosstalk through series of intersections. Crosstalk is measured on the first intersection of a series

2.6.2   Intersections

Unlike electronics, which uses multiple interconnection layers, waveguide intersections are unavoidable in many planar waveguide circuits. As any deviation from the perfect straight waveguide, an intersection introduces a scattering loss. Moreover, light can couple from one waveguide to the other either directly or indirectly through the weak resonance at the crossing point, resulting in crosstalk. Due to the large index contrast, SOI photonic wires have a high numerical aperture (NA) and thus the mode incident on an intersection will be highly divergent and a wide range of angles can be captured by the crossing waveguide. Also, there is a vertical mismatch between the waveguides and the crossing point. Simple intersections in Si wires were calculated by Fukazawa et al. [46] to have 1.4 dB loss and a crosstalk up to -9 dB. The resonance at the crossing point can be altered to reduce coupling to the crossing waveguide, which was studied by Johnson et al [47]. However, this reduces the transmission bandwidth of the intersection. By tapering the waveguides to broader waveguides, the mode incident on the intersection becomes less divergent, and one can expect direct coupling to be reduced. Also, the vertical mismatch is reduced, lowering reflections and scattering to radiative modes.

Orthogonal, four-fold symmetric intersections were characterised by measuring series of 2, 4 and 6 intersections in transmission. The crosstalk was measured on the first intersection. Structures with different taper widths and taper lengths were measured. Figure 2.18 shows the result with 1.5 μm taper width and 5 μm taper length, showing less than 0.65 dB loss and a crosstalk lower than -30 dB.

Figure 2.19: SEM picture of a waveguide intersection with double etch in order to reduce losses and crosstalk. The intersection has a length L, deep waveguide width w_d, and shallow ridge width w_s.

While a crosstalk figure of -30 dB will be acceptable for many applications, still only a few intersections in series can be supported due to the insertion loss. One possibility to further enhance transmission is lowering the index contrast. This was done here by using a double etch scheme (see Chapter 3). A SEM picture is given in figure 2.19. With a total length $L = 6\,\mu m$, and waveguide widths $w_d = 2.5\,\mu m$ and $w_s = 800\,nm$, a loss of 0.16 dB is achieved, and a crosstalk lower than -40 dB.

Through extensive simulation with 3-D software, it should be possible to further reduce the loss. For complex circuits, an order of magnitude improvement should be targeted.

2.6.3 Power splitters

Power splitters are essential components in any integrated circuit. For power division circuits and in wavelength-dependent filters, balanced 3 dB splitters are needed. In the other limit, a very asymmetric splitting ratio is desired for power taps for signal monitoring. A third function is as taps in wavelength filters with a parallel design, where a very fine and exact control over the splitting ratio is needed.

A 1×2 MMI 3 dB power splitter with a footprint of $3 \times 7.6\mu m^2$ was characterised through splitter trees with splitting ratios up to 1/64 (6 levels). The layout of such a splitter tree is shown in figure 2.20(a). The transmission for the different outputs is shown in figure 2.20(b). The measurement contained a lot of fluctuations. Fourier analysis of the spectra showed contributions from various cavity lengths, with strong compo-

nents near DC, corresponding to lengths between $10\,\mu m$ and $30\,\mu m$, and a component around $94\,\mu m$. The latter can either be due to the cavity between fibre and fibre coupler, or, less probably, to reflections at the ends of a $100\,\mu m$ length connection wire. The shorter cavity lengths may correspond to reflections in the splitters (MMIs) themselves and between splitters at either ends of short interconnection wires. Of course all the higher order reflections are present too, with exponentially decreasing magnitude. All these reflections generate spectral fluctuations of up to $1.5\,dB$, but we believe this is mainly due to the fibre coupling. A parabolic fitting was made to each transmission curve in order to avoid the noise, leading to curves with slightly detuned maxima, which is due to slight rotation of the sample with respect to the fibres and alignment differences. From these curves corresponding to 2 to 6 splitters, the transmission per splitter was extracted, shown in figure 2.20(c). The transmission shows a residual curvature, which is originating from the fibre couplers and not from the wavelength dependency of the splitters. The excess loss per splitter port is 0.15dB maximum. Around the optimum fibre coupler transmission wavelength, the excess loss is only $0.04\pm0.03\,dB$ per port, or $0.08\pm0.06\,dB$ per splitter. Also, from the fitted transmission curves the unbalance can be extracted. The unbalance between output ports is $0.3\,dB$ maximum, and clearly mainly due to extra bends and wire lengths between the outputs and not to the splitters, unless multiple unbalanced splitters would compensate for each other.

Next to the MMI structures, several Y-splitter types were characterised, but insertion losses are still around $1.5\,dB$. Further optimisation is necessary both in the theoretical design and the mask design.

2.6.4   Directional couplers

Directional coupling structures are essential for the ring resonators and Mach-Zehnder-based wavelength filters studied in chapters 5 and 6. However, while a very basic structure, the proper functioning of SOI directional couplers can be hindered by multiple factors. Scattering losses due to sidewall roughness are different for the symmetric and asymmetric supermodes, giving rise to a phase difference between the outputs different from $\pi/2$. Secondly, the sharp transition region at the input does not excite the symmetric and asymmetric supermodes in the proper balanced way. Both situations are equivalent to an asymmetric directional coupler, with a maximum transferred power less than 100%. A third problem is small width variations that actually make the coupler asymmetric. Fourth, the TE-like mode experiences some loss at interface between the straight and

(a) mask layout

(b) measured fibre to fibre transmission with MMI splitters

(c) loss of MMI splitter

Figure 2.20: Splitter trees for splitter loss and balance measurement. The fibre to fibre transmission through the tree for a MMI splitter is shown, and its fitted loss/splitter.
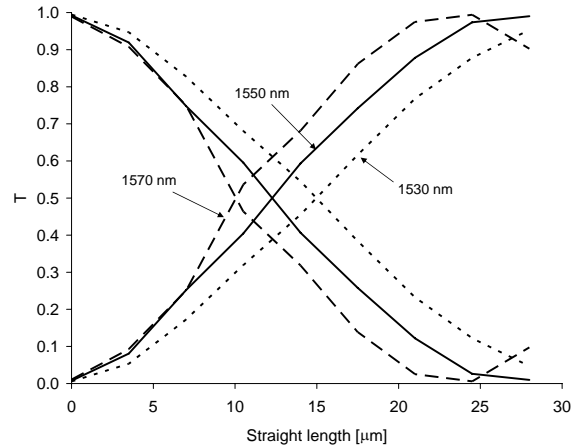
Figure 2.21: Measured normalised transmission of directional couplers as a function of the straight length. Bend radius is 6.2 μm. Measured waveguide width is 400 nm, gap is 270 nm. BOX thickness 2 μm.

bent waveguides due to mode mismatch. As simulation indicates that this power is mainly coupled to the TM-like mode, polarisation dependency problems may arise if the TM mode has a low enough propagation loss. Otherwise, power coupled to the TM mode is lost.

Figure 2.21 shows a measured transmission through both outputs of a directional coupler, as a function of the coupler length. The bend radius was 6.25 μm. Almost 100% coupling is achieved, with an extinction ratio better than 20 dB at the coupling length. For a zero straight length, with only bends coupled, 1% of power is coupled over.

2.7　Conclusion

In this chapter, we have studied the basic propagation properties of photonic wires, and we have seen a great sensitivity to dimensional variations. The phase velocity is varying greatly for dimensional variations of just 1 nm. This will have an important impact on wavelength filtering components, as will be shown in the next chapters. Nanometre-scale accuracy will be needed. Because of this sensitivity, we talk about nanophotonic waveguides and structures. Experimental data on basic circuit components has been discussed, and we see that although the loss and crosstalk performance look very promising already, a performance improvement of

more than an order of magnitude should be reached for really complex integrated functionality. Therefore, a more parallel integration scheme, integrating many short circuits on a single chip for parallel data handling, is more realistic in the relatively near future than the *city of light* on a chip envisaged by some researchers. Despite the nanometre-scale accuracy needed, structures can be made much smaller than in lower index-contrast material systems.

Chapter 3

Technology

High-index contrast, highly integrated photonic integrated circuits need a suitable fabrication technology. The most important drive to use silicon as a substrate material, is exactly the use of processes and tools developed for electronic IC fabrication. Within our research group, CMOS fabrication based processes were developed for fabrication of photonic IC's [1]. These wafer-scale processes were used in this work to fabricate photonic-wire based wavelength filters. In this chapter, after some remarks with respect to CMOS compatibility, an overview of the fabrication processes is given. Secondly, the possible sources of non-uniformity and inaccuracy are reviewed. In Chapter 8, these will be related experimentally to the numerical figures given in Chapter 2.

3.1   Silicon and CMOS compatibility

The huge knowledge and technology base, maturity and process control of CMOS fabrication processes are very attractive to be used for integrated photonics devices in or based on silicon. However, actual compatibility of photonics components with CMOS is a complex matter, which can be looked at on several levels.

First, one can use industrial tools that were developed for CMOS fabrication, such as optical lithography and etching machines, and materials and processes compatible with those tools, but use them in an environment not doing CMOS fabrication. The materials might not even be silicon, oxide or aluminium. This allows for the largest customisation of research, but is the least close to actual CMOS compatibility.

On a second level, one can actually use those industrial CMOS fabrication tools in an environment where CMOS fabrication is done. At this level, some boundary conditions appear:

49

- Contamination control is very strict in a CMOS environment. The list of materials allowed is typically very short. In a research environment, it is however possible, but not trivial, that new directions can be explored, such as processing III-V materials in an environment where silicon is processed.

- Related to the contamination control but also to process temperature, it is not straightforward to use a given kind of resist or process. For instance, if a certain resist or process is used in the CMOS back-end, it is not evidently possible to use the same material or process in the front-end. The flow of wafers can also go only from the front-end to the back-end and not vice-versa.

- Part or all of the process operations are done by operators that do not (necessarily) have a notion of integrated photonics. As the requirements for photonics and electronics are different, the photonic circuit designer or process developer needs to have a good understanding of the electronics language and way of seeing, to bridge the gap.

In a research environment, the processes used can be either plain CMOS processes, with limited use for general photonics, or adapted CMOS processes. In a commercial environment, one has to use the processes that are available.

On a third level, it is possible to be compatible with CMOS at a level that electronics can actually be integrated with photonic circuits on the same chip. This means that photonic structures are available as standard building blocks in design libraries, that the thermal budgets are fully compatible with the CMOS flow, and other requirements. This is a very hard level to reach for general silicon photonic integrated circuits, if one takes into account all the wishes of the photonic component researchers.

This work is situated at the second level, where both CMOS-compatible materials and CMOS tools are used, and fabrication is done in a CMOS pilot line, with only little interference from our part. The processes are adapted CMOS processes, optimised for fabrication of a large class of photonic structures. Mainly, this means trading the possibility to fabricate a large variety of structures for accuracy. As we will see, this is also the reason why the current fabrication process is not yet delivering the level of accuracy needed discussed in chapter 2.

Of industrial CMOS processes, the lithography process is probably the most limiting with respect to controllability of feature size and shape. However, it is one of the key tools that enables the high throughput and

high yield for electronic circuits. The number of academic research groups and research institutions actually doing deep UV lithography based processing for silicon photonics is limited. In Europe, these are IMEC [2] and LETI (Grenoble, France) [34], however fabricating designs for collaborating groups over Europe. Additionally, St-Andrews University (Fife, Scotland) has performed some tests in collaboration with Intel [48]. Outside Europe, an MIT research group [33] and a Chinese Academy of Science (Bejing) research group [49] have used DUV lithography. Two examples of the use of CMOS tools in a commercial environment are the startup Luxtera Inc. (CA, US), doing fabless SOI processing through a CMOS foundry and actual integration with CMOS, and Lightsmith Technologies Inc. (OR,US) using deep UV lithography to fabricate large gratings for diffractive optics.

3.2 Fabrication processes and fabricated structures

As the processes have been developed in the context of previous work [1], only a brief overview of them will be given here. Process development work was done with people from the Silicon Process and Device Technology (SPDT) division in IMEC-Leuven. An overview of fabricated components is given in Appendix B.

3.2.1 Process flow overview

Structures are created in the top Si layer of a 200mm SOI wafer by pattering a resist with 248 nm deep UV (DUV) optical lithography. The whole 220 nm thick layer is etched, creating a large horizontal refractive index contrast. For the creation of fibre couplers and other shallow structures, these steps can be preceded by another lithography step and more shallow etch step. All these steps happen at wafer scale in a semi-automated way. They can be fully performed within the Si pilot line at IMEC. After etch (and optionally after lithography), the structures are checked or extensively characterised using wafer-scale top-down SEM, and then diced manually or more recently by sawing. If waveguide facets of optical quality are needed, the substrate is mechanically thinned down to 200 to 250 μm before cleaving.

The thickness of the buried oxide was 1 μm for most fabricated structures in this work, and 2 μm for structures fabricated in the second half of 2006 (some AWG devices in chapter 7 and the directional couplers in chapter 2).

3.2.2 Basic optical projection lithography properties

The lithographical process used is optical projection lithography, with a wavelength in the deep UV range (248 nm). The demagnification of the optical projection system is 4X. The mask is a simple chrome layer, on a 4X reticle. The technology corresponds to the $.18\,\mu m$ CMOS node.

The field of optical lithography is very complex, and as of today, still not all the details of the physical and chemical processes during lithography, resist treatment and etching are well understood. In the following, some basic concepts are given, to give an understanding of the limits of optical lithography for integrated photonics. The interested reader is pointed to the excellent textbook by H.J. Levinson [50].

Resolution

The resolution of the lithographical process is basically defined by four parameters: the illumination wavelength λ, the numerical aperture of the imaging system NA, the *spatial coherency factor σ*, describing the coherence of the light source, and a material and processing dependent factor k. This is further explained in the following paragraphs.

The quality of the reconstruction of the mask pattern in the *aerial image* (the projection of the mask pattern in the resist layer when replaced with $n = 1$ (air)) depends highly on the structure type. For isolated structures, with a point as the illumination source, the smaller λ and the larger the NA, the more light diffracted from the mask pattern is captured by the lens system, and the better the reconstructed image. The achievable linewidth is roughly limited to $w > \frac{\lambda}{2NA}$. Degradation of the resulting image is gradual.

For periodic structures on the other hand, the lower λ and the higher the NA, the more diffraction orders are captured, and thus the better the aerial image. However, as light is diffracted into discrete orders, degradation is discrete. At least two diffraction orders are needed, so the pitch a must be larger than $\frac{\lambda}{NA}$. However, when considering a real source with a finite aperture, the situation gets better. Such a source is only partially spatially coherent. In the plane where the NA of the imaging system is defined, the image of the source is not a point anymore, but also spread out. For a periodic structure with a pitch smaller than $\frac{\lambda}{NA}$, the first diffraction orders can now be partially captured by the imaging system, and imaging is possible. This however comes at the expense of a decreased contrast, putting a tighter requirement on the resist non-linearity. The coherence can be changed by varying the size and form of the source. For peri-

odic structures, off-axis illumination, for instance by dipole or quadrupole sources, can be used in order to enhance resolution. The factor σ describing the level of spatial coherence, is defined as the ratio of the radius of the image of the illumination source in the NA plane, and the NA, for simple illumination types. For a perfectly coherent source, $\sigma = 0$. The resolution of the aerial image of a periodic pattern is thus limited as

$$a > \frac{\lambda}{NA\,(1 + \sigma)}.$$ (3.1)

Therefore, to enhance the resolution of periodic patterns, a large enough σ is used, which however decreases the contrast for isolated features.

The influence of the resist, including its contrast, and the processing is described by a factor $k > 1$. Ideally, the resist has a step-like response. However, a practical resist has a less perfect non-linear response. The resist development and etch have an imperfect behaviour that can be modelled in this factor k. Therefore the practical resolution limit is larger than given by eq. 3.1.

For fabrication of photonic structures, the resolution matters mainly for the resolution of fine features. The main features are of optical wavelength scale, of the order of hundred to a few hundred nanometres. For today's state of the art optical lithography, these are large. However, for nanophotonic waveguides and structures, the accuracy of sub-wavelength features matters, and therefore the resist properties have a large influence.

Process windows

The position and slope of the transition in the developed resist profile, defining a feature, depends on the aerial image profile and the non-linear resist response. The exact feature width therefore depends on the exposure dose. Also, within the resist layer, the vertical profile at a given point should be as constant as possible, so the imaging system should have a large enough depth of focus (DOF) for the targeted feature pitches. However, the larger the NA, the smaller the DOF of the imaging system. Within a certain focus range (DOF) and within a certain exposure energy range, the *exposure latitude*, the feature with a given pitch will be considered to be on target. [1]. Therefore, for a given critical dimension (CD)[2], a *process*

---

[1]Note that moving to a higher NA can increase the fraction of the first diffraction orders transmitted for large enough pitch, and thus increase the resulting DOF of the full lithography process for larger pitches, while for smaller pitches the resulting DOF decreases

[2]The critical dimension (CD) on a CMOS layer is the smallest or most important linewidth, trench width or hole diameter on that layer
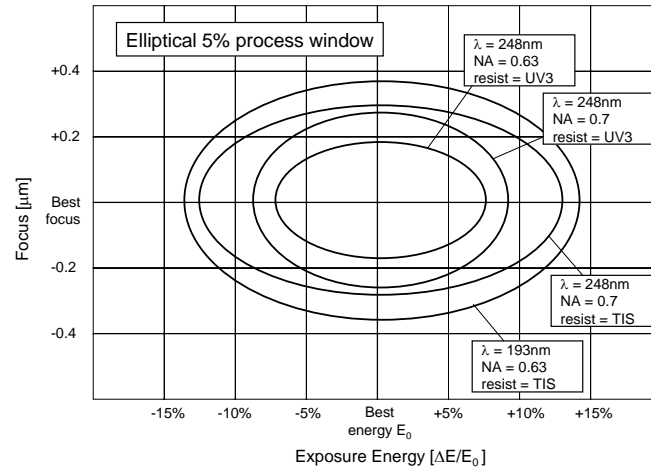
Figure 3.1: Process windows for a triangular lattice of holes with a modest 500 nm pitch and 300 nm diameter, for 4 different lithography processes. Reproduced from [1]

.

window can be defined by the depth of focus and the *exposure latitude*. In CMOS, it is common to work with a 10% allowable deviation, as designs are tolerant to such deviations or can easily be made so. However, as we will see in the following chapters, this is a much too relaxed specification for wavelength-selective photonic devices. An example of such a process window is given in figure 3.1. It is elegantly described by the largest ellipse that can be inscribed in the actual measured 5% deviation contour. Process windows for multiple combinations of illumination wavelengths, resists and NA are given.

For nanophotonic structures to be fabricated well, the process windows should be extremely large compared to a typical CMOS process. Typically, a deviation of 1% is already very large. Enhancing the process windows will enable structures with different pitches to be on target at the same time, which is needed for nanophotonics. The windows will also have an influence on how well the sidewall slope, sidewall width deviations within a feature and sidewall roughness can be controlled. The details of the resists are of extreme importance in this respect.

As was already mentioned in the previous paragraph, isolated and periodic structures behave rather differently. Through control of the NA and the coherence of the light source, periodic structures can be more control-

lable then isolated features. This has an influence on the process windows, and on the reproducibility of structures. Therefore, adding assist features to isolated structures, structures on the mask near the structure that do not show up in the developed resist, may help to improve on their reproducibility.

Resist properties

The resists used for DUV lithography processing are chemically amplified. The illumination releases an acid which diffuses during a post-illumination bake step and then catalyses a deprotection reaction rendering the illuminated resists soluble (positive tone resist) or insoluble (negative tone resist). There are a multitude of resist types using different reactions, such as cross-linking for negative tone, and deprotection or depolymerisation for positive tone [51].

Relevant to this work is that the diffusion range of the acidic catalyst has a large but complex influence on both the sidewall roughness of the resist pattern and on the resolution (included in the k factor). Optimising both the sidewall slope of the resist, and a possible slope of the etched structure, and the sidewall roughness and other deviations at the same time proves a difficult task. For a single feature type, for instance lines with 450 nm width, this is possible. The problem lies in controlling these aspects while combining lines with periodic structures of different kinds and sub-wavelength structures.

Reflection problems

If the mask is projected in a resist layer directly applied to the substrate, reflection problems can arise. Light is reflected both at the air-resist and resist-substrate interface, resulting in complex reflection phenomena depending on the substrate topology and the resist thickness and properties. Two major effects arise: standing waves in the resist, and a variation of parameters and the CD with the resist thickness, called swing effects. For instance, the *dose-to-clear* (the absorbed dose needed for making the resist soluble), can vary with resist thickness, resulting in such CD variation. The swing effects can be particularly severe for DUV lithography, even with limited topography. The resist thickness should be chosen at an extreme of the swing curves, so that the influence on the process window is smallest. Several reflection reduction strategies exist. The two major solutions are application of a top or bottom antireflective coating (TAR or BARC, respectively). A TAR does not avoid the reflection from the sub-

strate, and therefore standing waves can still result. However, swing effects are reduced in amplitude, so the deterioration of the process window is reduced. A BARC strongly reduced or avoids the substrate reflection, and the resulting standing wave patterns, but is more difficult to optimise than a process with a TAR.

3.2.3   Process flow with 220nm etch

The process starts from a bare SOI wafer with a 220 nm thick silicon top layer and a 1 μm buffer oxide. The Si substrate has a thickness of approximately 750 μm. An 800 nm thick layer of Shipley UV3 resist is spun on top of the wafer. The resist is a standard positive tone resist for contact holes with 248 nm DUV lithography[3]. All illuminated features are therefore etched away. The NA used is 0.63, and the spatial coherency factor $\sigma = 0.4$. After a pre-bake of the resist, a top anti-reflective (TAR) coating is applied. The resist is then patterned by an exposure with 248 nm DUV lithography, a post-exposure bake and development.

The process was developed to support the fabrication of dense periodic arrays of holes, photonic crystal structures, as well as isolated features. Therefore, a *dark-field* mask[4] is used in combination with positive tone resist. Strip waveguides are therefore defined by only etching away trenches and not all of the material except the waveguide.

The silicon top layer is then etched using an ICP-RIE process with a $Cl_2/He/HBr/O_2$ plasma, with the patterned resist as the etch mask. Afterwards, the residue of the resist is removed.

For the second generation of components (PICCO02 PWIRE, see Appendix B), the patterned resist was directly used as an etch mask. The propagation loss of waveguides fabricated with this process was reported in chapter 2. Losses as low as 2.4 dB/cm were obtained.

However, this process proved sufficient for isolated line structures such as photonic wires, but not for periodic structures of varying pitch and fill factor. The combination of both is even more difficult. The main problem was the much too large bias in feature size between lithography and etch for periodic structures, and the difference in bias between isolated and periodic structures. For all later generations, the resist was first given a plasma treatment, hardening and shrinking the resist, compensating for a

---

[3]The main reason for this is that the process was developed for fabrication of photonic crystals.

[4]A dark field mask has an opaque *field* and structures are defined by transparent features on the mask. With a positive tone resist, the transparent features on the mask will therefore be etched away.

(a) titled view   (b) cross-section

Figure 3.2: SEM pictures of photonic wire with standard process

large litho/etch bias at the cost of the lithography process window. The propagation loss of straight wire waveguides obtained with this second process is about 0.5 to 1 dB/cm higher than with the first process. Nevertheless, the second process was used throughout this work.

3.2.4 Fabricated photonic wires

Figure 3.2(a) shows an SEM picture of a wire bend fabricated with the standard process. One can see there is still a lot of sidewall roughness. Visually, we can determine that this roughness is curtain-shaped with a characteristic pitch of the order of 50 nm. An actual measurement of this would need high quality SEM pictures and advanced top-down image analysis methods or atomic force microscopy (AFM). Additionally, the top of the waveguide seems to be etched through at the borders, resulting in small holes. This is due to resist breakthrough, a result of the plasma treatment. SEM pictures of photonic wires with the process without plasma treatment did not show the microholes on top. The influence of the breakthrough on the sidewall roughness is unclear. While this plasma treatment enabled the fabrication of photonic crystal structures and the combination with photonic wires, it thus introduced a higher scattering loss. The measured loss values are consistently around 3 to 3.5 dB/cm.

The cross-section in figure 3.2(b) shows the top silicon is etched for about 40 nm each side, the sidewall slope is 9° and there is a slight footing of 15 nm width and 15 nm height, which is most probably slightly etched into the oxide layer.

Figure 3.3: Top-down SEM images of photonic wires. The wire edges contain very small nm-scale fluctuations on a length scale of a few 100 nm along the wire. The measurement bars given on the right picture are only approximate.

Along the propagation direction, the sidewall width is slightly varying. From figure 3.3, short length scale roughness cannot be seen, but there is a certain *waviness* of the waveguide, with width fluctuations on a length scale of some hundreds of nm. These sidewall fluctuations are in the range of several nm, which will induce a lot of phase errors as discussed in chapter 2. SEM pictures of waveguides fabricated without a plasma treatment given to the resist showed similar variations, indicating that the resist hardening did not have much influence on or cause these longer length scale variations.

To obtain more accurate information on these sidewall variations on a longer lengths scale, top-down SEM images were further analysed. Figure 3.4 shows an example, with a wire of almost 1.7 µm long. It should be noted that the result will depend on the length of wire measured. First, the wire edges are detected using a simple threshold algorithm on gray-scaled images. The sidewall deviation is obtained for each sidewall by subtracting its mean. Histogram of the deviations are plotted in figure 3.4(b). The Fourier transforms of both sidewalls are given if figure 3.4(c). These figures show a clear contribution of deviations in the 0.6 to $10/\mu m$ range (with $0.6/\mu m$ the smallest measurable frequency given the limited length of the wire), or periodicities in the 100 nm to 1.7 µm range.

3.2.5 Two-level processing

For the creation of fibre couplers (Section 2.3.2) and other shallow structures such as waveguide crossings (Section 2.6.2) and star couplers (Chapter 7), the 220 nm deep Si etch can be preceded by a 50 nm or 70 nm deep Si etch. The alignment of lithography steps is done on wafer-scale using standard markers with an accuracy of 50 nm. The structures used here can be designed so that this level of accuracy introduces no excess losses.

If the wafer can stay on the stage in the stepper, the alignment accuracy is better, at about 6 nm. Therefore, multiple lithographies could also be used for structures with the same etch depth without too much interface losses for a large class of structures. However, this comes at the cost of added reticle real estate. Figure 3.5 gives some examples of structures fabricated with double etching: a grating coupler and a star coupler from an AWG. From the half star coupler picture, we can see the alignment between the two lithography steps is very good.

(a) SEM image

(b) Histogram of deviation

(c) Fourier transform of deviation
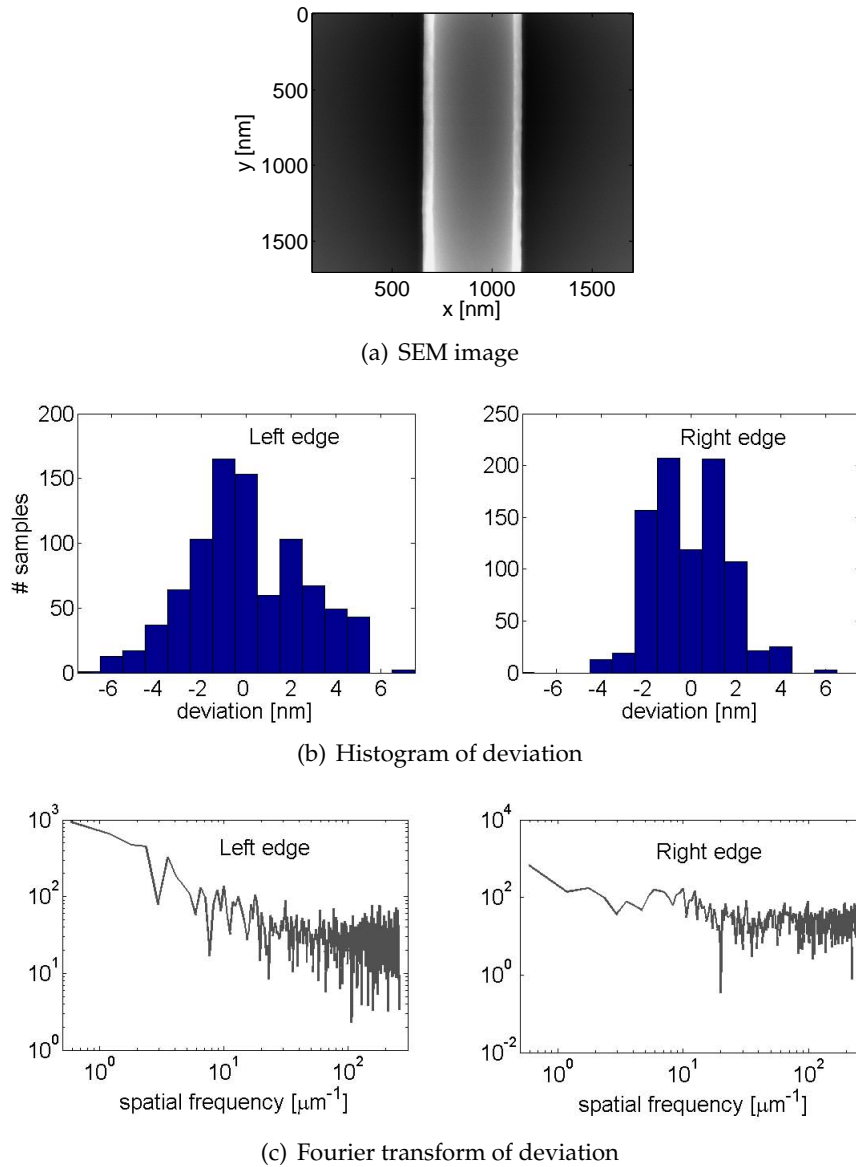
Figure 3.4: Analysis of the sidewall deviations. The SEM image is first grayscaled and normalized on its minimum intensity. The sidewalls are traced along the y axis using a threshold of 0.6 times the maximum intensity of the normalized image. The mean position is subtracted from the resulting coordinates. Histograms of both sidewall deviations are plotted in (b). The Fourier transforms (spatial spectra) of the deviations are plotted in (c).

(a) fibre coupler (b) fibre coupler (c) input part of star coupler

Figure 3.5: SEM pictures of structures with two-level processing

3.3 Alternative fabrication methods

Using CMOS tools for prototyping of photonics devices is expensive. The fabrication methods are optimised for high volume fabrication, dividing the cost over a high number of components. For research, the cost gets divided over a very small number of components instead. Also, the resolution of the optical lithography tools is limited to some tens to hundreds of nm, which is good enough for a lot of nanophotonic devices, but not for the general class of structures being researched. A third reason why CMOS fabrication is not optimal for research, is that one wants to fabricate a large range of structures with very different feature sizes in a single fabrication step, so that no alignment problems arise. However, CMOS processes are optimised for a single feature type, such as quasi-periodic arrays of gate lines or contact holes. Therefore, the processes must be adapted and compromises must be made in order to fabricate photonic circuits. A fourth reason is the design to fabrication cycle can be long, as a lot of designs need to be grouped on the expensive masks in order to be affordable.

Other fabrication methods are used for research on nanophotonic circuits, that deliver high resolution and short fabrication cycles.

- Electron beam lithography (e-beam, EBL) uses a serial writing process in resist, with a resolution depending on the beam diameter. Structures can be written with very high resolution compared to DUV lithography, and being a serial process, every structure can be tuned on an individual basis. There are proximity effects between structures, but the effect is additive and thus relatively easy to model. The main disadvantage is the slow nature of the serial writ-

ing process, making it unsuitable for volume fabrication and dense integrated structures. Also, the write field of an e-beam writer is limited to some hundreds of micrometre, before the wafer stage needs to be translated, causing stitching errors. Nevertheless, as a single structure can be written in a relatively short time, and resolution is high, it is the preferred tool for research.

- Focused ion beam lithography (FIB) is similar to EBL in operation, but uses a focused Ga^+ ion beam that is scanned over the material. The beam can be used both for observation (at low energies) and fabrication. Direct milling into the silicon is possible, but an etch mask process can be used too. Compared to EBL, the FIB tool is more versatile, at a comparable resolution. However, for fabrication of photonic components, it is still in the early stages of research. The main bottleneck is the understanding and reduction of the damage done by the ions during the etching process.

- Nano-imprint lithography is a replication technique which is, in contrast to EBL and FIB, suitable for volume fabrication, and is not as fundamentally limited in resolution as DUV lithography. Additionally, as no expensive optics are needed, it is in principle a very cheap process. Fabrication of a 1X mold is needed, which can for instance be made through EBL. Several techniques for transferring this mold in the material exist. For instance, a transparent mold can be used, and after imprinting in resist, the resist can be fixed with a UV exposure through the mold. Fabrication of photonic structures is still rather in the early stages of development, but can be expected to advance fast.

3.4 Sources of non-uniformity and inaccuracy

Every step in the fabrication process is to some extent a source for unintended variations. Here, the most influential variations and errors are listed, together with indicative values for their magnitude with the fabrication process used for all components in the following chapters. From the simulation results on the sensitivity of wire structures in section 2.5, combined with the simulation results, we can get an idea of the fabrication process tolerances and what is needed to achieve high reproducibility, accuracy and uniformity. This is discussed in chapter 8, based on measurements of the wavelength selective components studied.

3.4.1 Mask digitisation errors

All structure coordinates on masks are snapped to a 5 nm by 5 nm grid (1X coordinates) by the mask definition software. Care is taken so that no unintended 5 nm width differences and offsets are present in the resulting mask design. Most structures are defined as polygons, described by their vertices. However, the chrome pattern on the actual reticle depends on how the intermediate file formats and software handles the drawing of the polygons. Mask digitisation can differ from the original file 5 nm grid, and thus sub-5 nm differences can exist between structures on the reticle. Also, line drawing algorithms will create a best fit line between vertices, but there may be nm-scale variations between oblique lines. The same can happen in bends (that are actually defined as polygons). Therefore, care must be taken with structures with lines not along the raster lines and bends in different directions.

3.4.2 Sidewall roughness

Fabricated waveguide sidewalls are typically vertically patterned with a curtain-shaped roughness. This is clearly visible in figure 3.2(a), where a roughness profile with an amplitude of a few nm and a typical period of a few tens of nm can be seen. This roughness has multiple possible sources and can be described with stochastic models.

Already on the mask, there is a certain degree of line edge roughness (LER). Recent measurements on 4X reticles for recent technology nodes have shown a roughness with a 3σ of around 7 nm with a correlation length of the order of 500 nm [52] (values on mask). However, when measuring on larger length windows, another autocorrelation became apparent, in the range of several µm [52]. On the wafer, this mask LER would give rise to roughness frequency components around 125 nm and 500 nm.

Secondly, all three of the lithography, resist trimming and etching steps contribute to the roughness. The patterned resist typically already contains line edge roughness (LER). The sources of LER are not yet completely understood, but the magnitude of LER in these kinds of resists depends on many factors. Amongst these are resist characteristics such as molecule size and the molecular weight of the polymer and the acid diffusion range in the resist [53, 54]. The aerial image contrast can also have an influence [55]. A third source of influences is the development step [56]. The patterns present in the resist are then transferred to the silicon by the etching process, either enhanced or smoothed.

Another source of sidewall and top wall roughness in the current process can clearly be seen on figure 3.2(a). Microholes are formed in the top wall near the sidewalls. Most probably these are formed near the end of the etching process, when the resist is totally consumed at the sidewall. As this effect was not seen in wires processed without the resist hardening step and the hardening step slightly trims the resist, we can assume this resist breakthrough is related to that plasma step. These holes will lead to additional scattering, and can at the very end of the process influence the etching of the sidewalls and thus give rise to additional sidewall roughness.

3.4.3   Mean critical dimension variation

There are a number of sources of deviation of the critical dimension (CD) from the designed value: variations on the reticle, the existence of multiple CDs on a single mask, illumination dose variations, optical proximity effects and corner rounding, and roughness and other stochastic width variations within a line. Some of these are elaborated on here.

Variations on the reticle

Variations on the reticle are translated into variations on the wafer. As a 4X projection system is used, the absolute variations on the reticle appear reduced in the projected image. However, imaging of the mask features is not linear, and the change of resist feature size per unit mask feature size change (divided by 4) is described by a mask error (enhancement) factor (MEF or MEEF). Generally, the MEF>1 and the relative error in the resist is thus larger than the relative error on the reticle.

The CD variation that is already present on the reticle is measured by the mask shop through top-down inspection. For the structures fabricated here, 0.18 µm technology is used, with a specified tolerance on the deviation of the mean CD from the nominal CD, and with a specified uniformity of the CDs about their mean. For a 720 nm line/space structure, both specs are ± 22 nm. This translates in a ≥ 5.5 nm spec on the wafer for a 180 nm line/space, or a $\geq 3\%$ tolerance.

As an example, for the PICSOI08 reticle, measured mean deviations from the nominal CD were -12 nm and -9.5 nm for two separate sets of CD measurements, one in the X and one in the Y direction. On the wafer, this translates into at least -3 nm and -2.4 nm, respectively. The maximum deviations from the mean were +5.1/-3.9 nm (+1.3/-0.98 nm) and +6.4/-4.2 nm (+1.6/-1.05 nm) for the two sets. For contact holes, the same

specs hold, and measured CD variation was similar. Measurements on the other reticles showed similar values. In conclusion, while the reticles are within spec, for photonics applications we can expect nm-scale deviations from the designed value, and variations between lines of the order of ± 1 to ± 3 nm, just due to linewidth variations on the mask.

Influence of the illumination conditions

For CMOS, a mask layer contains structures with a very limited range of pitch and feature sizes. For fabrication of photonic structures with DUV lithography, the situation is totally different, as wildly different structures are combined on a single layer in order to avoid alignment mismatches. It is clear from figure 3.1 getting a variety of structures on target at the same illumination dose. For different structures, the process windows will at best have an overlapping region. One can then choose to work at the optimum dose for the most critical dimension, or choose a dose in the overlapping region. Depending on the exposure dose latitude, the CDs will then differ from their designed value.

As an example, we take the experimental data from [1] (figure 5.19) on isolated lines. In order to print a 500 nm wide line on target, the optimal exposure dose at optimal focus is[5] about $22\,mJ/cm^2$. The optimal dose for a 700 nm wide line is about $24\,mJ/cm^2$. This 10% difference in dose would give a somewhat less than 10% difference on either CD, if the lithography was optimised for one of them. It is clear that for nanophotonic circuits, this is a huge difference.

It is possible to compensate for optimal dose differences in the mask design, by simple application of a CD bias. This needs extensive process characterisation, which was done in [1]. Taking the same example, at the optimal dose for a 500 nm isolated line, a 700 nm wide line should be designed 680 nm wide. However, when optical proximity effects are added to this (see further), and a larger variation of feature sizes, such corrections become cumbersome or impossible to get perfect.

The further the illumination dose is from the optimal dose, the less tolerant the feature is to small focus and dose variations, that could be present between illumination of different wafers (and in theory also between dies).

---

[5]It should be noted that after a recalibration of the 248 nm DUV stepper in 2006, the dose-to-target for 500 nm wide isolated lines is $18mJ/cm^2$.

(a) top-down (b) cross-section

Figure 3.6: SEM pictures of two wires in close proximity (≈200nm gap)

Optical proximity effects

In optical lithography, the beams diffracted from the structures in the mask pattern will interfere. As this can happen constructively or destructively, the influence of densely packed structures on each other will show a complex behaviour. These effects are called optical proximity effects (OPE). One notable example in the context of this work is the linewidth difference between an isolated (*iso*) wire and two wires that are in close proximity such as in a directional coupler. An example of this proximity effect with the current process is given in Figure 3.6(a). One clearly sees that the wires in close proximity are narrower than the isolated wire, and the gap is wider at least when seen from the top. From a cross-section view of a similar structure (Fig. 3.6(b)), we can see that the slope of the inner sidewalls is different from the outer sidewalls. In such line pairs, the gap width is limited by the resolution. With the current process, gaps down to 150 nm have been fabricated, however by underexposing the structures: using an illumination dose below the dose to target for an isolated waveguide, therefore broadening the waveguides a lot.

Of course these proximity effects can have a severe impact on the optical behaviour. By applying *optical proximity corrections* (OPC) on the mask, OPE can be compensated for in the actually fabricated structures. OPC were applied (PICMOS02) to line pairs, by broadening the lines in the pair compared to the isolated lines and shrinking the gap, but using the dose-

Figure 3.7: Isolated lines and line pairs with and without OPC measured after etch, for different linewidths and 180nm gap on the mask

to-target for isolated lines. Figure 3.4.3 shows the measurement result, where each wire was made 10 nm broader towards the gap and thus the gap was 160 nm on mask instead of 180 nm. After lithography, the width of the lines in the line pair is 5 to 10 nm off the isolated line width, while the gap is 180 nm to 190 nm. The same line pair without OPC has gaps between 210 nm and 230 nm, and linewidths that are 30 nm narrower than the isolated lines.

The presence and severity of optical proximity effects will very much depend on the details of the lithographic process and the wavelength used. It has been shown for instance, that with 193 nm DUV lithography it is possible that no visible optical proximity effects appear for photonic crystal waveguides [48], while these clearly show up in our current process.

Corner rounding, line end shortening and tapered trenches

At the end of a line or trench, the line/trench is typically shorter than designed. This can be corrected for by making the line slightly larger on mask or using assist features such as *hammerheads* (a local widening of the line end). A related effect is the rounding of line corners and filling of trench corners. This is illustrated in figure 3.8. A better resolution will

(a) intersection (not tapered)   (b) Y splitter

Figure 3.8: Examples of corner rounding

reduce these effects, but nevertheless they should be taken into account during design.

3.4.4   Etch depth variation

The etch rate is dependent on the fill factor of the resist on the mask. For photonic structures, the fill factor can be different between masks, especially masks for research purposes, as some will contain mostly isolated features such as wire structures, and others will contain very dense structures. Therefore, the etching depth may be slightly different between wafers with different masks.

Secondly, also within a wafer, slight etching depth variation may exist. For instance, etching depth variations of a few nm were measured using AFM on 45 to 50 nm deep etched fibre couplers, see [57](figure 6.11). Depending on the structure, such a variation may be important. For the grating couplers, we generally see small spectral shifts, impeding comparison between structures. From figure 3.2(b), we can see the buried oxide layer (BOX) is unintendedly etched slightly (about 10 nm). Therefore, small etching depth variations will have little influence on the modal effective index.

3.4.5   Thickness Variations and wafer defects

The wafers used have a specified mean thickness variation between wafers of 20nm 3σ. This means a standard deviation of the wafer-to-wafer wave-
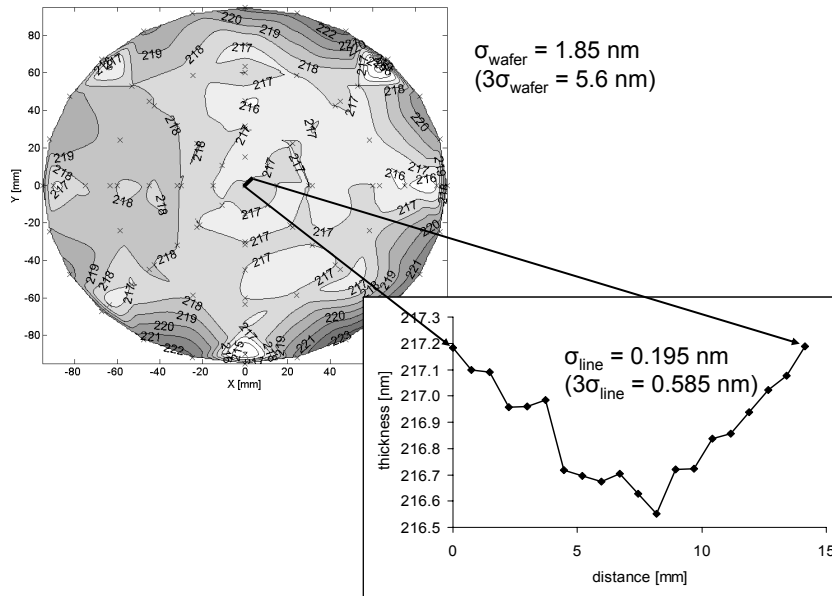
Figure 3.9: Measured silicon top layer thickness on a wafer with 220 nm Si, 1000 nm BOX nominally. Courtesy S.K. Selvaraja.

.

length shift of about 10nm (see chapter 4)! The wafer thickness also influences the etch depth to some extent.

Within a single wafer, the specified maximum non-uniformity is $\pm 13\,nm$. In 1997, rms values as low as 0.13 nm were measured [58], while the manufacturer's website has an example with an rms of 0.32 nm. This means a maximum wavelength non-uniformity of 21 nm and a rms shift in the 200 pm to 500 pm range. For WDM, this causes wavelength shifts of a considerable fraction of the channel spacing or larger. Colleague S.K. Selvaraja measured the thickness and refractive index distribution of the Si top layer over the wafer, for a blank standard SOI wafer used throughout this work. The result is given in figure 3.9. Over the wafer, the mean thickness is 218 nm, with a standard deviation of $\sigma = 1.85$ nm. However, the strongest variation is between the rim and the centre of the wafer. Taking all measurement points located at least 10 mm from the wafer edge, the mean thickness in the centre is 217 nm with $\sigma = 0.8$ nm. The edge points have a mean thickness of 220.8 nm with $\sigma = 1.32$ nm.

In some cases, localised defects such as microholes (HF defects) can be present in the wafer that are not directly visible by visual inspection but

can be revealed with photoluminescence inspection [59]. The wafers have a specified maximum of 0.3 HF defects / cm^2.

Through AFM measurements of the silicon surface of a wafer by colleague S. K. Selvaraja, we have also obtained an idea of the top surface roughness. The obtained values are used in in chapter 2 to calculate the scattering loss. The top surface roughness has an rms of 0.098nm within a $1\,\mu m \times 1\,\mu m$ measurement window, as shown in figure 3.10(a). The measurement was repeated for a different location on the wafer, and both locations were measured in a $1\,\mu m \times 1\,\mu m$ window, yielding very similar data. The autocorrelation of each scanline was taken, giving consistent data between scanlines. The autocorrelation of the data on a scan line may be approximated with an exponential function with a correlation length of about 100 nm, shown in figure 3.10(b). The top of the main lobe fits better to a Gaussian autocorrelation function with a correlation length of 30 nm.

3.4.6 Stress Variations

Due to the different thermal expansion coefficient and the relaxation of stress at 1100°C during wafer bonding in the Unibond process [60], stress is present in the different layers at room temperature. While partially relaxed by wafer bow, an in-plane compressive strain is present in the buried oxide, and a tensile strain is present in the top silicon and the upper part of the substrate. Measurements on SOI with a 205 nm top Si layer and a 400 nm BOX showed a limited stress of 2.3MPa in the top Si, a -290MPa stress in the BOX and a stress in the substrate that reaches a maximum of 235MPa at the Si/SiO$_2$ interface [61]. We can reason that if a 2.3MPa stress is induced by a temperature change of 1100°C, stress changes of only 2.1 kPa per °C temperature change can expected. While we can expect the stress not to be drastically different with a 1μm BOX, [61] shows the situation does change when a top oxide is added. Also, etching of the top Si pushes the stress up in the top layer, while the BOX relaxes. For photonic wires, the Si layer is completely etched through and a two-dimensional structure is created, which will alter the stress compared to a slab structure. Intuitively, we can expect the stress in the wire to rise and the compressive stress in the BOX to relax. Still, stress changes as a function of temperature will be limited. From the figures in section 2.5, we see the effect on wavelength dependent components is smaller than 1 pm and thus negligible. There will be stress differences between isolated lines and line pairs however.
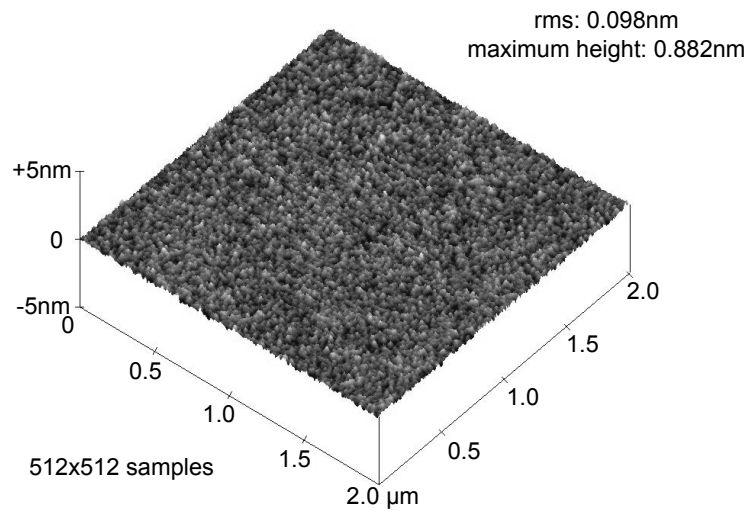
(a) AFM scan

(b) scanline at y=39 nm

Figure 3.10: AFM measurement of the silicon top surface on a wafer with 220 nm Si, 1000 nm BOX nominally. The measurement window was 2 μm × 2 μm, with a resolution of 3.9 nm both directions.

3.4.7   Improving the processes

How can the processes be improved ? On the lithography side, a lot can be done. The current process was initially developed for densely packed holes, not for isolated lines. Of course, in reality one will want to combine both these extremes and intermediate structures, putting strong requirements on the process.

Switching to a shorter illumination wavelength, 193 nm, will deliver a higher resolution. While this is not actually needed for many photonic devices, moving to 193 nm lithography, together with advanced resists, can also enhance the process windows and reduce proximity effects. The enhancement of the process windows can for instance be seen from figure 3.1. An absence of proximity effects was for instance seen in [48]. Choosing a resist with a response that is better suited for isolated lines will also improve the line quality, but may deteriorate the quality of dense structures and holes.

Using a thinner resist, will result in a smaller DOF requirement. However, a certain thickness is needed as resist is consumed during the etching. Therefore, a hard mask could be used for etching instead of resist. Such a hard mask can be much thinner than a resist mask, and thus the resist mask needed to pattern the hard mask can be thinner as well, and the required DOF smaller. The choice of a hard mask or resist mask is however not straightforward. With a suitable resist and a sufficient DOF, more straight sidewalls can be obtained. The sidewall roughness can be diminished by choosing a proper resist too, and through optimisation of the etching process, but there will be a complex trade-off between roughness and slope, for instance due to the acid diffusion range. In any case, the microholes on top of the current structures should be avoided.

Development of a more advanced process is currently the focus of colleague S.K. Selvaraja.

3.5   Conclusion

With the current process, small enough linewidths can be obtained that seem to be reasonably well controllable if precautions are taken during the design phase. Roughness and small longer scale width variations are present, and sidewalls have a 9° slope. The fabrication process was developed for photonic crystal structures and a deep etch through the buried oxide. It was later adapted for a shallow etch through the core only. Nevertheless, photonic wire waveguides, which the process was not optimized

for, show very low propagation loss as low as 2.4 dB/cm, as reported in chapter 2.

The process was further adapted to enable the fabrication of photonic crystal structures, and to combine these with photonic wire structures. The resulting process introduced microholes in the top surface of the isolated wire waveguides near the sidewalls, and therefore yields waveguides with a higher propagation loss, typically of around 3 to 3.5 dB/cm. Nevertheless, this process was used throughout the major part of this work,

A number of structures have been demonstrated with the current process and are discussed in chapters 5 to 7. For a higher integration level, higher performance with respect to loss and crosstalk, and for a higher accuracy and reproducibility, the process should be enhanced. As we have seen, the resolution of the process as such is not a real problem, at least not for wire bases structures, although with a higher resolution accuracy could be enhanced. A more important problem is that of sidewall variations, both on a short length scale (roughness) and on a longer scale (wavy waveguides). However, reducing these variations takes looking for the cause, but multiple sources are possible: the reticle, the resist behaviour during lithography, post-bake and development, and the etching. Already on the reticle, variations of the linewidth are present, translating in nm-scale variations on the wafer. This is beyond our control, as well as the quality and accuracy of the bare SOI wafers. Through control of the mask and resist loading, the global and local transmission factor of the mask and the relative resist coverage), through the use of assist features, and through further process development, a better control can be obtained in the future.

Chapter 4

Integrated optical filters

Before starting the discussion of specific wavelength filters in chapters 5 to 7, an introduction is given to the type of integrated optical filters studied in this work. Then, the filter specifications for optical wavelength filters for WDM are discussed. Not for all applications and for all requirements, a well-defined spec exists, however. Then, a basis of general filter design is given, mainly in relation to chapter 6, the interesting relation between optical an digital filters is reviewed. Finally, the basic parameters of WDM filters in silicon-on-insulator are given, and their sensitivity to fabrication variations is calculated, based on the figures reported in chapter 2.

4.1 Optical filters

All filters discussed are based on interference between single-mode light paths. In the case of a ring resonator (chapter 5), a light path interferes with itself. In a Mach-Zehnder interferometer (MZI) and in MZI cascades (chapter 6), two light paths interfere. Multiple light paths interfere in the case of lattice filters (chapter 6) and arrayed waveguide gratings (AWGs, chapter 7).

With such single-mode, interference-based structures, complex filters can be obtained. The design and optimisation of such filters is well known for low-index-contrast systems. However, direct implementation in the high-index-contrast SOI system is not straightforward.

Other approaches towards wavelength filtering are possible. Especially in photonic crystals, periodic structures that work on the basis of distributed omni-directional reflection, specific features of the spectral behaviour of the periodic and quasi-periodic structures may be exploited. One example is the wavelength-selective coupling between modes in a structure. For instance, coupling between two waveguides can be medi-

ated selectively through higher order modes [10]. A second example is exploiting the highly dispersive nature of these structures [11].

4.1.1   Filter types: high-level view

In an optical network based on multiple wavelengths carrying data, several wavelength-dependent functions are needed. A typical core or metropolitan network today that is based on wavelength division multiplexing, has a hybrid topology. For instance, it may consist of multiple interconnected rings, with star networks at each node. In a star topology or for the interconnection of multiple rings, wavelength routers and cross-connects are needed, spatially routing wavelengths between input and output fibres (figure 4.1(a)). Such a router can be fixed, supporting only a static routing configuration, or reconfigurable, allowing for much more flexibility. Fast reconfigurability may be useful for fast switchable networks. At each simple node in a ring, an add-drop filter is needed (optical add-drop multiplexer, OADM) as shown in figure 4.1(b), capable of dropping one or more wavelength channels downstream, and capable of adding data on the same channels upstream. Again, these filters may be static, or reconfigurable (ROADMs).

Other wavelength-selective devices are necessary as well, such as gain equalisation filters or dispersion compensation filters.

In this work, we have developed both channel drop filters (without explicitly studying the *channel add* functionality), and static wavelength routers.

4.1.2   Filter types: low-level view

The actual optical operation of a wavelength filter can be very diverse. As already mentioned, only filters based on interference between a single mode with itself or between multiple single-mode light paths were studied here. Three general classes of filters can be distinguished.

A first class of filters employs feed-forward paths only. Examples are the Mach-Zehnder and Michelson interferometers, both with a delay length in one path compared to the other path. These filters intrinsically have wide passbands, but it is not very easy to obtain a good stopband extinction. A second class of filters uses feedback paths. These are the resonant filters, such as the Fabry-Perot interferometer and the channel drop outputs of disk and ring resonators coupled to two waveguides. Resonant filters intrinsically have a small passband width, but a high stopband extinction. Finally, both feedforward and feedback paths may be combined,

(a) Wavelength router

(b) Add-drop multiplexer
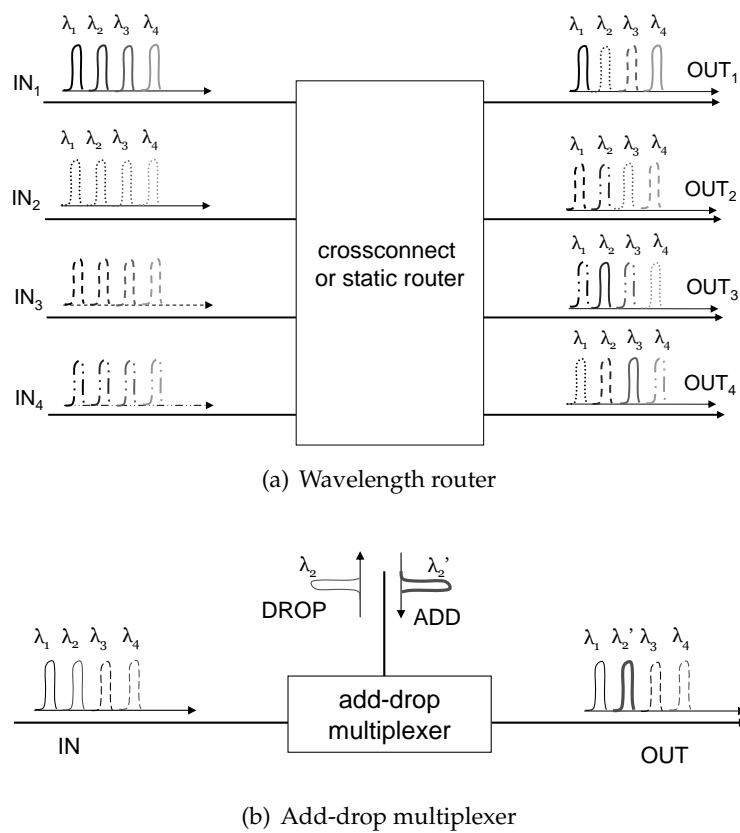
Figure 4.1: Schematic operation of an add-drop multiplexer and a wavelength router or cross-connect

with mixed features of both former types. In a single device, multiple outputs may even exhibit different behaviour resorting under one of the three classes, as is the case for the complementary outputs of a ring resonator channel drop filter.

Standing and travelling wave resonators

A difference between the Fabry-Perot and a ring or disk resonator is that the former is a standing wave resonator, while the latter is a travelling wave resonator. When used as channel drop filter for instance, this has important implications: a single standing wave resonator coupled to a waveguide couples power to both directions in the waveguide simultaneously. Therefore, a single standing wave resonator cannot transfer power to a single output with 100% efficiency, while this is possible with a travelling wave resonator. However, combining two standing wave resonators can alleviate this problem [62].

Higher order filters

Several basic filter stages (with a single delay length) may be combined in order to obtain better bandwidth and roll-off. For each basic stage type, there are multiple possibilities to combine stages. Mach-Zehnder stages may be cascaded in a tree-like structure, as in figure 4.2(b). The properties of each stage can then be tuned to obtain a desired filter shape. Another possibility is to cascade them in a quasi-periodic manner as in figure 4.2(a), in which case the filter is usually called a lattice filter. The whole parameter space of the filter should then be tuned at once to obtain a desired spectrum. Ring resonators can be cascaded into ring lattice filters too, as illustrated in figure 4.2(c). Again, the coupling strengths between the rings are optimised to approximate the desired transmission spectrum. By combining Mach-Zehnder stages with ring resonators, for instance as in figure 4.2(d), complex higher order filters can be obtained. The ring resonators are coupled to single waveguides and act as all-pass filters, giving a very non-linear phase shift around their resonance wavelength. By loading the arms of a Mach-Zehnder with multiple all-pass rings, complex behaviour is possible. In this work, Mach-Zehnder lattice filters and logarithmic Mach-Zehnder cascades are studied in chapter 6. The fabrication variations did not yet allow to successfully obtain ring lattice filters beyond second order, as reported in chapter 5.

The filter types above all have a serial nature and are all single or multiple channel drop filters. By using multiple such filters in parallel on a

(a) Mach-Zehnder lattice

(b) Mach-Zehnder cascade

(c) Ring lattice

(d) Combining resonant and non-resonant behaviour

Figure 4.2: Some filter types

single input, $1 \times N$ multiplexing or demultiplexing and $M \times N$ routing of many channels is possible. However, single-component optical filters exist that already have a very parallel nature and can act as multiplexers or routers. These are mainly the concave grating and the arrayed waveguide grating. The former is not discussed in this work, but is a very interesting type of filter for which the high index contrast has some advantages such as a smaller device size and better tolerance to sidewall slope compared to lower index contrast devices [63]. Silicon-on-insulator AWG devices have been designed and fabricated and will be discussed in chapter 7.

4.1.3 Filter specifications

The most important filter specifications for this work are the channel frequencies or wavelengths, the period of the spectrum (free spectral range), the required stopband or passband widths, the tolerated insertion loss, the tolerated loss non-uniformity, the required isolation between channels in a given output, and the required isolation between outputs for a given channel. These specifications are illustrated in figure 4.1.3, based on a fabricated and characterised arrayed waveguide grating device (chapter 7). Additionally, the polarisation dependence (polarisation dependent loss and polarisation dependent spectral shift) and temperature dependence are important specifications for general filters. The temperature dependence of SOI filters will be studied in chapter 8 and in chapter 9, a polarisation independent AWG-based component is shown.

WDM filters for core and metro networks

Components for use in the core of telecommunication networks typically have specifications that cannot be met yet with nanophotonic SOI technology. Nevertheless, it is interesting to list the specifications of such components, as telecommunication technology can also be used for shorter range communication.

Wavelength channels The specifications for WDM telecommunication applications are formalised by the ITU [1]. The DWDM (dense WDM) and CWDM (coarse WDM) grids are defined in standards G.694.1 and G.694.2, respectively. A de facto standard on many component specs was established by Telcordia [2] with recommendations on most communication technology. The ITU specifies an equidistant frequency grid, with a channel

---

[1] International Telecommunication Union - www.itu.com
[2] Telcordia Technologies, Inc. - www.telcordia.com

Figure 4.3: Filter specifications, illustrated with a measurement example of an eight-channel arrayed waveguide grating in SOI with a 400 GHz channel spacing, equidistant in frequency.

spacing Δf_{ch} of 12.5 to 200 GHz for DWDM, and channel frequencies of 193100 GHz+$n \cdot \Delta f_{ch}$. For metropolitan networks or access networks, often uncooled lasers are used in combination with *coarse* WDM (CWDM) with channel spacings of 20 nm.

Additionally, most components such as widely tuneable lasers operate in the *C-band*, between 191 THz (1570 nm) and 196 THz (1530 nm).

Insertion loss and channel isolation The loss and channel isolation (crosstalk) requirements vary greatly between applications, and are not strictly specified. In telecommunication, isolation requirements are generally higher than for datacommunication and interconnects, as the power budget is normally smaller. A typical good telecom receiver has a sensitivity of -20 dBm or better to reach a BER of 10[-9] at 10 Gbit/s, and the modulated output power of a telecom source is typically about 1 mW. This means a power budget of about 20 dB. Without amplification, a lot of this will be consumed by propagation and connector losses, splitter losses, wavelength filter losses and other loss sources. With many channels propagating through the system together, this leaves little room for crosstalk. For WDM and DWDM, commercial component isolation specs are typically

25 dB for the adjacent channels and 35 dB down to 50 dB for non-adjacent channels. For CWDM, used in metro networks, typical isolation specs are 20 to 30 dB.

Bandwidth   In telecommunication, lasers with very narrow linewidth and temperature control are used. Therefore, the required bandwidth is basically defined by the modulation speed and format. With NRZ coding, this means a practical bandwidth requirement of approximately 10 or 40 GHz for 10 or 40 Gbps communication, respectively. For metro applications, often uncooled lasers are used, and therefore the required bandwidth is mostly defined by the allowed source wavelength tolerance.

Access network filters

Optical access networks typically operate with specifications differing a lot from those of core and metro networks. Currently, fibre-to-the-home systems typically use different bands around 1300 nm and 1550 nm wavelength for downstream and upstream communication. Sometimes, different services are put on different bands or subbands. At the access point, duplexers or triplexers are needed to separate the widely spaced bands, and uncooled lasers are used as sources. Specifications on loss and isolation are more relaxed than for telecom applications. This enables cost-effective, consumer-ready modules.

By employing WDM techniques and technology, a greater flexibility and greater aggregate bandwidths can be obtained. For instance, one can use telecom lasers and a dense grid of wavelengths in the C-band, to support multiple users over a single passive optical network (PON), with flexibility in allocation of channels and bandwidth. Due to the splitting losses in the PON, the loss and isolation requirements are still tight. A component for such an application is discussed in chapter 9.

WDM techniques in datacommunication

For data communication between modules, electronic boards, or even processors, typically shorter wavelengths are used, with VCSEL sources and multimode fibre. However, it is not a bad idea to use the high performing telecom components for specific applications, such as high bit rate, reconfigurable interconnects. For these applications, communication speeds will be as high as those in telecom networks, and therefore loss and isolation specs will be rather strict. Losses other than wavelength filter losses will be smaller than in a telecom communication path, but amplification

may be impossible. A component for such an application is discussed in chapter 9.

4.2 Filter concepts and design

In this section, some basic concepts of filters are discussed, their high-level design procedures and issues, and the relation to digital filters.

All filters studied here are based on interference between light paths with a delay relative to each other. Multiple filter stages can be concatenated in series, or filters can have a number of parallel paths. Therefore, the filters can be described based on a unit delay length (or delay time), where every path length differs from the other path lengths by an integer multiple of the unit delay length. One can deviate from whole multiples of this unit delay length in order implement complex functionality or to design structures that are more tolerant to fabrication variations.

4.2.1 Unit delay and free spectral range

The unit delay length (time) of a filter determines the periodicity of its transfer functions, the free spectral range (FSR). The unit delay time is related to both the physical length and the group velocity. Especially in high-index contrast material systems it is important to take the group index into account and not the effective index of the waveguide mode[3].

For proper operation, the width of a pulse travelling through a filter should be significantly longer than the unit delay length, as otherwise aliasing can result. In the frequency domain, this means the bandwidth of the pulse should be significantly smaller than the FSR.

The delay length can either be set from the required FSR Δf_{FSR} or from the required absolute channel frequency f_0:

$$\Delta f_{FSR} = \frac{c}{n_g \Delta L} \tag{4.1}$$

$$f_0 = \frac{mc}{n_{eff} \Delta L} \tag{4.2}$$

with n_{eff} the effective index, n_g the group index, ΔL the unit delay length and m an integer (the filter order). If the delay lines are composed of sections with varying n_g, such as bends and straight wires or wires of differ-

---

[3]In many textbooks and papers one finds the effective index, which is a valid approximation for glass-based low index contrast systems with small waveguide and material dispersion.

ent width, obtaining the right FSR becomes more cumbersome or limits the possible designs.

For a demultiplexing device, the central frequency and the channel spacing are design parameters, and the delay length is set according to the central frequency and a sufficiently large FSR. Other applications such as an AWG wavelength router may require a cyclic frequency behaviour of the outputs and thus need a well controlled FSR. The delay length will then be set from the required FSR. If $n_{eff} \neq n_g$, the right frequencies will need to be obtained through tuning. The tuning method should of course mainly affect n_{eff} and have a small impact on n_g.

4.2.2 Relation to digital filters

With each lightpath having a length that is an integer multiple of a fixed unit delay length, and filters containing feed-forward and/or feedback paths, there is a clear correspondence between discrete-time, digital filters and continuous-time, analogue optical filters. Even the actual conceptual functioning and structure of optical filters with discrete time delays and digital filters have a lot in common. Therefore, it is no wonder that all the filter design methods of the digital filter world can be pulled in to design optical filters in a structured way. Of course, a mapping between the parameters governing the optical filter functioning and those of digital filters needs to be made: the magnitude and phase part of coupling and splitting ratios in the optical domain need to be related to the complex coefficients of the polynomial descriptions or the pole and zero locations in the Z domain from the digital filter world. An excellent introduction to this was written by Madsen and Zhao [64]. In chapter 6 such methods are actually used for filter design.

For an optical filter based on interference between delay lines, a unit delay length ΔL introduces a phase factor $\mathrm{e}^{-j\beta\Delta L}$, with β the propagation constant, or $\mathrm{e}^{-j\omega T(\omega)}$, with T the unit delay. A linear, time-invariant (LTI) filter is fully described by its impulse response $h(nT)$, for $n = -\infty \ldots \infty$7, which is a discrete series. The transfer function is the discrete time Fourier transform of the impulse response, but the frequency domain can be analytically extended to the complex z-plane. The transfer function is then given by the Z-transform of the impulse response as

$$H(z) = \sum_{n=-\inf}^{+\inf} h(n) z^{-n} \qquad (4.3)$$

Having a discrete-time impulse response, the transfer function of the optical filter is periodic, with a free spectral range (FSR) $\Delta f_{FSR} = 1/T$. Equation 4.3 is then evaluated on the unit circle to translate to the frequency (and optical frequency) domain, by $z = e^{j\beta\Delta L} = e^{j\omega T}$. For an optical filter with dispersive waveguides, T is given to the first order as $T = \frac{n_g\Delta L}{c}$.

Poles and zeros

Each filter can be written as a function of its poles and zeros in the z plane. This is exactly where optical filter design through its relation with digital filters can show its power. The pole and zero locations can be analysed and optimised through well-known techniques from the discrete-time filter world. The transfer function of a general filter can be written as

$$T = Cz^{(N_p - N_z)} \frac{\prod_{k=0}^{N_z}(z - z_k)}{\prod_{m=0}^{N_p}(z - p_m)}, \tag{4.4}$$

with C a normalisation factor.

4.2.3 Impulse response

Discrete-time filters with feed-forward paths only have a finite impulse response (FIR). These filters are also called moving average (MA) filters The length of the impulse response and the step response are equal to the filter order, or the number of possible paths with a different total delay length. Therefore, a non-resonant optical filter with discrete delay lengths, such as a Mach-Zehnder lattice or cascade has a finite impulse response.

On the other hand, digital filters with feedback paths only have an infinite impulse response (IIR), and are sometimes called autoregressive (AR) filters. Filters with feed-forward and feedback paths have an infinite impulse response as well, but their transmission also has FIR filter features. Therefore, a better name for these is autoregressive moving average (ARMA) filter. A filter may exhibit transmissions to certain outputs of the AR type and to other outputs of an ARMA type. In the optical domain, such is the case for a thin-film filter.

4.2.4 Magnitude and phase response, delay and dispersion

A zero on the unit circle translates into zero transmission in the optical domain at that frequency. Zeros inside or outside the unit circle generate a dip in transmission with a depth that is inversely related to the distance

of the zero from the unit circle. The magnitude response for zeros locates symmetrically around the unit circle is equal.

Poles, optically relating to a resonant cavity, translate into a maximum in transmission in the optical domain. As passive optical filters are always stable, their poles should be located within the unit circle. With gain, unstable filters (lasers) can be obtained, and poles are located outside the unit circle. The off-resonance extinction of a single-pole filter is higher as the pole approaches the unit circle.

The group delay of a filter is given by the slope of the phase response of the filter:

$$\tau_g = -T\frac{d}{d\omega}\arctan\frac{\Im H(z)}{\Re H(z)}. \tag{4.5}$$

Around a resonance, sharp transitions of the phase transfer can result, and the delay can lose its meaning, and even become negative, for instance in the pass port of a ring resonator add-drop filter (see chapter 5).

The group delay due to a single zero is negative for a zero with $|z| < 1$, a minimum phase zero. A zero outside the unit circle is a maximum phase zero, and its group delay is positive. The group delay increases as the zero approaches the unit circle, but the group delay of a linear phase zero ($|z| = 1$) is $0.5T$. Multiple zeros contribute group delay additively. Note that a single-stage optical FIR filter can be minimum phase due to causality, unless the group delay loses its meaning due to the physical process (strong distortion of the pulse). Due to the stability requirement for optical filters, a pole always contributes positive delay, which increases as the pole approaches the unit circle.

Causality also dictates that it is impossible to obtain an infinitely sharp roll-off and wideband zero transmission in the stopband, as this would require an impulse response extending to $n = -\infty$.

The filter dispersion is defined as the derivative of the group delay with respect to frequency. With the $\tau_n = \frac{\tau_g}{T}$ the normalised group delay and $\nu = fT$ the normalised frequency, the normalised dispersion is $D_n = \frac{d\tau_n}{d\nu}$. For an optical filter, the dispersion in ps/nm is then given by

$$D = -c\frac{T^2}{\lambda^2}D_n. \tag{4.6}$$

4.2.5 Linear phase filters

Filters with zeros or poles close to the unit circle, will exhibit a large dispersion. For a FIR filter, if all zeros are inside the unit circle, a minimum phase filter results and the group delay will be smallest. A filter with all

zeros outside the unit circle is a maximum phase filter, and the group delay will be maximum. The filter's group delay can become negative and lose its meaning around the frequencies of zero transmission. An optical filter with poles only is always minimum-phase, due to causality and stability.

A linear phase filter, exhibiting a flat group delay, is advantageous for optical filtering as it does not distort the pulse shape. As an AR filter is always minimum-phase, only FIR filters can exhibit linear phase, if their zeros are properly positioned. Two obtain a linear phase filter, all zeros should appear on the unit circle, or in reciprocal pairs around the unit circle. One can prove that the impulse response of a linear phase filter has a Hermitian symmetry:

$$h(n) = (-1)^P e^{j\Phi_P} h^* (N - n),\qquad(4.7)$$

with P the number of roots on the unit circle, and Φ_P the sum of the phases of these roots.

A linear phase filter with real coefficients has an even or odd symmetry of the coefficients (and thus the impulse response), depending on the filter order. The transfer function has the same symmetry, making the magnitude response always symmetric around the centre frequency.

Optical bandpass filters are designed as low-pass filters, given their high carrier frequency which makes the FSR much smaller than their central frequency. Such a filter can only be linear-phase if the symmetry of the coefficients is even, independent of the filter order.

Channel drop filters dropping multiple bands generally have an asymmetric magnitude response around a centre frequency, and therefore must have complex coefficients. In the optical domain, this means that in general phase shifters need to be applied in each filter stage. A filter with complex coefficients can be designed to be linear phase, if its coefficients adhere to equation 4.7. Other applications such as gain equalisation need filters with complex coefficients as well.

4.2.6   Filter design

Three basic approaches for designing optical filters are possible.

The most accurate but most time-consuming option is to simulate the whole structure with an electromagnetic simulation tool such as eigen-mode expansion (EME), beam propagation (BPM) or finite difference time domain (FDTD) method based software. Of course, this is only a good approach for a limited set of filters. With FDTD, only extremely small filters can be simulated. With BPM and EME, the filter should function almost entirely along a single propagation direction.

Secondly, one can simulate the structure on a higher level, changing coupling coefficients, propagation constants, phase shifts and loss and optimising the calculated filter response towards a desired transfer function. Parts of the structure can be simulated with the appropriate tools in order to obtain the right properties, for example the coupled power in a directional coupler. However, such optimisation may take a long time if the starting values are not properly chosen.

The third approach, but not suitable for all optical filter types, is first synthesise and optimise a filter on a high level with one of the numerous methods available for design of digital or analog filters, and translate the obtained filter coefficients (under the form of equation 4.4) to the optical domain properties of the filter structure. This method is extremely well suited for Mach-Zehnder and ring lattice filters, or for example for ARMA filters based on all-pass ring resonators and MZI structures. The difficulty is in translating between the digital and optical domain. For Mach-Zehnder lattice filters, this can be done based on the work by Jinguji and co-workers [65]. This approach was used in chapter 6.

Digital FIR filter design

For design of digital lowpass FIR filters, several methods are available. One can start from a sinc as an impulse response (with a rectangular transfer function), and apply a window function to reduce the Gibbs ripple around the transition band at the cost of a worse roll-off (larger transition band). For instance, a Hamming window can be used. A better roll-off may be obtained using a least-squares or minimax error minimisation. The error is optimised over the specified passbands and stopbands. With the least-squares method, the integrated squared error over the different bands is minimised. With the minimax method, the maximum error is minimised (Chebyshev approximation). This latter method is interesting as it produces equiripple designs. The Parks-McClellan algorithm was developed to implement a minimax method [66]. This method produces the steepest roll-off for FIR filters, and was used in chapter 6 in conjunction with the Jinguji method to translate the filter coefficients to coupler coefficients in a Mach-Zehnder lattice filter. An extension of the Parks-McClellan algorithm allows for designing filters with complex coefficients. This was not done in this work.

Digital IIR filter design

Methods for design of general ARMA-class digital filters are numerous. For instance, Butterworth filters have maximal flat passbands and stopbands, but suboptimal roll-off. Chebyshev type I and II filters show a steeper roll-off at the same filter order, but have an equiripple behaviour in the passband or stopband, respectively. Elliptic filters have the steepest roll-off, but show equiripple in both passband and stopbands. Obtaining a large enough extinction with the latter filter thus needs the largest filter order.

In this work, only three simple ARMA-class filters were put on a mask for testing: a second-order and a third order ring lattice filter, and a third order Buttherworth filter consisting of a Mach-Zehnder with an all-pass ring resonator in one arm. These will be discussed at the end of chapter 5. The main problem with filters employing multiple rings is the inaccuracy between ring resonators, even if in close proximity on the mask and wafer.

4.3 WDM filters in silicon-on-insulator

4.3.1 Waveguide dispersion and unit delay length

Due to the large difference between the group index and the effective index in a high index contrast waveguide, it is often impossible to adhere to the ITU WDM specifications without tuning: either the FSR or the central frequency deviate from their desired values. The ITU specifies channel spacings of 12.5 GHz, 25 GHz, 50 GHz and so on, and channel frequencies of 193.1 THz plus an integer multiple of the channel spacing. In a low index contrast system, $n_{eff} \approx n_g$, and $\Delta f_{FSR} = \frac{f_0}{m}$, with m the filter order. If ΔL is set from the required FSR, as for a cyclic routing device or an interleaver, then f_0 exactly coincides with one of the WDM channels for all possible channels. However, for a general system, $\Delta f_{FSR} = \frac{f_0}{m} \frac{n_{eff}}{n_g}$. When ΔL is set from the required FSR, then f_0 never coincides with a WDM channel. Vice versa, if ΔL is set from f_0, then the FSR is never an integer multiple of the channel spacing. Other applications, such as wavelength multiplexed systems for metro and access networks, may of course have other specifications.

Table 4.1 gives an idea of the delay lengths needed when calculated from the FSR, based on 450 nm wide, oxide clad waveguides. The unit delay time $\Delta \tau$ is also given. The filters in chapters 6 and 7 are all designed with a delay length computed from the required FSR.

Δf_{FSR} [GHz]	$\Delta\lambda_{FSR}$ [nm]	$\Delta\tau$ [ps]	ΔL [µm]
100	0.8	10	661.79
200	1.61	5	330.90
400	3.22	2.5	165.45
800	6.43	1.25	82.72
1600	12.86	0.625	41.36
3200	25.73	0.313	20.68
4800	38.59	0.208	13.79
6400	51.46	0.156	10.34

Table 4.1: Delay lengths needed to obtain a given FSR, around 1550 nm wavelength, for a waveguide width of 450 nm and oxide clad waveguides.

If the delay length is set to obtain a given central frequency, then the FSR is not exact. For a component such as a demultiplexing device based for instance on an AWG with many wavelength channels, this is not a problem.

4.3.2 Sensitivity to variations

For general filters based on interference, either resonant or non-resonant, the centre of the passband or stopband will be given by

$$\lambda_f^0 = n_{eff}\left(\lambda_f^0\right)\frac{\Delta L^0}{m} \tag{4.8}$$

with ΔL the physical delay length of the filter (such as arm length difference in a Mach-Zehnder, ring circumference in a ring resonator), m the filter order and n_{eff} the mean effective index of the waveguides out of which the filter delay length is composed. In a simple model, a single waveguide index can be considered. In reality, designed waveguide width variations, directional couplers and other structure influence the index. However, the simple model may still give a good indication of spectral changes due to variations of either n_{eff} or ΔL. When calculating deviations of λ_f, the large waveguide dispersion will need to be taken into account. In general,

the variation of λ_f around λ_f^0 due to some effect z is given by

$$
\frac{d\lambda_f}{dz} = \left(\frac{\partial n_{eff}}{\partial z} + \frac{\partial n_{eff}}{\partial \lambda_f} \frac{d\lambda_f}{dz} \right) \frac{\Delta L^0}{m} + \frac{n_{eff}^0}{m} \frac{d\Delta L}{dz}
$$

$$
\Rightarrow \frac{d\lambda_f}{dz} = \frac{\frac{\partial n_{eff}}{\partial z} \frac{\Delta L^0}{m} + \frac{n_{eff}^0}{m} \frac{d\Delta L}{dz}}{1 - \frac{\partial n_{eff}}{\partial \lambda_f} \frac{\Delta L^0}{m}}
$$

$$
= \frac{n_{eff}^0}{n_g^0} \cdot \left(\frac{\partial n_{eff}}{\partial z} \frac{\lambda_0}{n_{eff}^0} + \frac{\lambda_0}{\Delta L^0} \frac{d\Delta L}{dz} \right) \qquad (4.9)
$$

The first factor in eq. 4.9 corrects for the waveguide dispersion. It is therefore important to note that, if $\frac{d\Delta L}{dz} = 0$, the shift is inversely proportional to the group index, not the effective index. For oxide clad wires the correction factor varies from 0.52 to 0.65 for widths between 400 nm and 600 nm, respectively. A more detailed plot of this factor is given in figure 4.4.

The second term in eq. 4.9 gives the sensitivity to a cavity length variation. If the cavity length variation with the influence parameter is absolute, then this term is inversely proportional to the cavity length (or cavity order), and thus filters with a longer free spectral range will exhibit a larger spectral shift. The term becomes very important for cavity lengths that approach the wavelength. For instance, this holds if the cavity length would be changing due to a lithographic effect. However, if cavity length variation scales with the cavity length, such as for linear thermal expansion, then the second term is *not* inversely proportional to the cavity length. Variations of the global waveguide width, height or refractive index, give a shift that is independent of the order of the filter, for a given filter wavelength. Of course, when comparing different orders within a single filter, the shift *does* depend on the order.

Height variations   The correction factor is larger for wider waveguides, making the wavelength shift grow. The wavelength shift for a typical filter is 1.3 to 1.4 nm shift per nm waveguide height difference. More detailed data are given by figure 4.5(a).

Width variations   The shift due to width variations is given in figure 4.5(b). As the change in refractive index becomes much smaller with increasing width, the wavelength shift decreases with wire width even though the correction factor grows: at 450 nm width, 1 nm width variation gives a shift of about 1 nm, while at 800 nm, the shift is only 200 pm.

Figure 4.4: n_{eff}/n_g (shift correction factor) at 1550 nm

(a) height variation

(b) width variation

Figure 4.5: Wavelength shift due to height and width variations, as a function of waveguide width. $\lambda = 1.55\,\mu m$.

Figure 4.6: Wavelength shift due to core and top cladding refractive index variations, as a function of waveguide width. $\lambda = 1.55\,\mu m$.

Refractive index A change of the core refractive index of 10^{-3} induces a wavelength shift of the order of $0.4\,\text{nm}$. This is depicted in figure 4.6. The cladding index has a smaller influence, but for small wires, the sensitivity can reach $100\,\text{pm}/10^{-3}\text{RIU}$ for e.g. water solution claddings. This large shift opens up possibilities for chemical sensing.

Temperature The refractive index change gives a wavelength shift of 68 to $73\,pm/K$. In order to predict the total temperature-dependent wavelength shift of filters, the dependency of refractive index on stress and the thermal expansion need to be taken into account. The thermal expansion coefficient $\alpha = \frac{1}{L}dL/dT$ of Si is $2.5 \cdot 10^{-6}/K$, while the expansion coefficient for silica is an order of magnitude smaller. If the expansion of the core was not hindered by the cladding, the wavelength shift is given by eq. 4.10.

$$\frac{d\lambda_f}{dT} = \frac{\frac{\lambda_f^0}{n_{eff}^0}\frac{\partial n_{eff}}{\partial T} + \lambda_f^0 \alpha_{Si}}{1 - \frac{\lambda_0}{n_{eff}}\frac{\partial n_{eff}}{\partial \lambda}} \qquad (4.10)$$

In this case, the expansion of the wire cross-section would only be a few pm, which is unnoticeable. In the direction of propagation, the effect is a very small shift of $4\,\text{pm/K}$. In reality, the real effect will be smaller as the silica cladding forces a tensile stress in the core that relaxes with rising temperature.

In conclusion, the temperature dependent wavelength shift can be expected to be in the 70 to 80 pm/K range. Interesting to note is that this

shift is irrespective of the cavity length (filter order, FSR). This large shift is attractive for temperature sensing.

Non-linear effects   In filters based on resonant structures, non-linear effects can have an influence on the refractive index through carrier effects and secondary thermal effects. A change of refractive index induces a spectral shift, which is in this case depending on the input power and the actual cavity details. With SOI wires, these effects can already play at modest power levels [24].

Phase variations and other effects   Stochastic or deterministic width and height variations not only induce a wavelength shift, but also generate phase errors after propagation through a delay line. Directional couplers, splitters and other structures suffer phase and power transmission variations too. When this happens within a single filter, this leads to sidelobes in the filter spectra, reduction of transmission or extinction in the passband or the stopband and a reduction of the roll-off. These effects will be studied in the chapters 5 to 8.

4.4   Conclusion

Wavelength channel filters in silicon-on-insulator form an interesting research subject, as they are a basic component in telecommunication networks and are very sensitive to the fabrication process quality. For more short-distance communication, the component specifications are often more relaxed. Therefore, nanophotonic SOI components may find their use first in such applications, if polarisation and temperature control are not needed or can not be used. Optical filters with delay lines with lengths that are a discrete multiple of a unit length, relate closely to digital filters. Therefore, the entire set of digital filter design methods may be applied to them. This will be used in chapters 5 and 6.

Chapter 5

Ring resonators

The ring resonator uses optical feedback, letting a travelling wave interfere with itself. Coupling the resonator to two waveguides, depicted in figure 5.1(a), leads to a transmission spectrum from the *input* to the *drop* port with a Lorentzian-like lineshape. In case of a lossless cavity, all power can be extracted from the input waveguide at resonance and transferred to the drop port. Therefore, such a configuration acts as a channel drop filter. By exciting the *add* port with a signal at the resonance wavelength, the dropped signal can be replaced and one obtains an add-drop filter.

The main advantage of using a travelling wave resonator is that all power can be coupled to a single output, whereas a single standing wave resonator coupled to two waveguides always couples to multiple output ports[1]. Compared to filters with only feed-forward paths, the ring resonator intrinsically offers a small bandwidth. To obtaining a wider bandwidth, a higher-order filter consisting of series coupled resonators is needed.

Another possibility for using the ring resonator is the all-pass configuration of a ring coupled to a single output, as shown in figure 5.1(b). In case of a lossless filter, the bus waveguide has unity power transfer but a steep phase step in the phase response around resonance. This all-pass filter can be used as part of a more complex filter, using the resonator as a complex phase rotator. If the cavity experiences loss, a Lorentzian lineshape appears in the power transmission.

Apart from its wavelength filter function, the ring resonator can be used to implement switching and modulation by changing the resonance frequency or extinguishing the resonance by introducing cavity loss.

In this work, we focused on single ring resonator filters. While these have limited usage, their successful fabrication with CMOS fabrication

---

[1]Combining two standing wave resonators can solve this problem [62].

(a) channel drop filter

(b) all-pass filter

Figure 5.1: Basic filters based on a single ring resonator

methods is a first step towards more complex ring resonator based filters. Single ring filters and a couple of higher order filters were experimentally tested to demonstrate the viability of SOI and wafer-scale technology for ring resonator based devices. In section 5.1, single-ring channel drop filters will be discussed. Calculations were done based on a high-level circuit model and realistic material parameters such as propagation loss. Single-ring resonators filters were fabricated and characterized. Section 5.2 discusses single-ring all-pass filters, that give a non-linear phase shift around their resonance. Further on, a simple ring-resonator based demultiplexer is shown. Finally, in section 5.5, two experimental examples of higher-order filters based on ring resonators are shown.

5.1 Single ring channel drop filters

First a brief overview is given of the transmission equations governing the ring resonator for CW excitation. These can be used to design ring resonators, not only statically, but partly also dynamically by looking at the transmission bandwidth. In the next section, we will pay a little more attention to the dynamic behaviour.

The drop characteristic of a single ring can easily be described as the equivalent of a digital single-pole filter with an infinite impulse response (IIR). The through port transmission has both a pole and a zero. For a passive structure, the maximum pole magnitude is unity and the filter is always stable. The zero is a minimum phase or on the unit circle in the z-plane for a lossless ring.

(a) lossless channel drop filter   (b) lossy channel drop filter

Figure 5.2: Power and phase transmission spectra of a single ringresonator channel drop filter with and without loss. Full line:trough port transmission, dash-dotted line: drop port transmission

5.1.1 Background

Steady-state model

Let κ be the amplitude coupling between ring and bus waveguides, and $\tau = \sqrt{1 - \kappa^2}$ in case of a lossless coupler. Let A be the amplitude transmission of half a roundtrip around the ring, excluding coupling to the bus waveguides. The definition of the ring radius can be chosen rather arbitrarily, but the effective index $n_{eff} = n_{eff}^R + j \cdot n_{eff}^I$ is always related to the radius definition chosen. With α the amplitude extinction coefficient of the travelling wave in the ring (in $1/cm$), and L half the ring roundtrip length, $A = \mathrm{e}^{-\alpha L} = \mathrm{e}^{-\frac{2\pi}{\lambda} n_{eff}^I L}$. Let $\phi = 2\beta L$ be the phase accumulated over a roundtrip.

Transmission   For a symmetrically coupled resonator, the drop port power transmission is, after some calculation,

$$T_d = \frac{T_{max}}{1 + F \sin^2\left(\frac{\phi}{2}\right)} \tag{5.1}$$

$$T_{max} = \frac{\kappa^4 A^2}{\left(1 - \tau^2 A^2\right)^2} \tag{5.2}$$

$$F = \frac{4\tau^2 A^2}{\left(1 - \tau^2 A^2\right)^2} \tag{5.3}$$

The pass (through) port power transmission is

$$T_p = \frac{r_{min} + F \sin^2(\phi/2)}{1 + F \sin^2(\phi/2)} \tag{5.4}$$

$$r_{min} = \frac{\left(1 - A^2\right)^2 \tau^2}{\left(1 - \tau^2 A^2\right)^2} \tag{5.5}$$

The extinction ratio at resonance between drop and through port is given by $E_r = \frac{T_{max}}{r_{min}}$, the maximum off-resonance extinction ratio between through and drop port $E_{ro} = \frac{r_{min} + F}{T_{max}}$.

Free spectral range   Within a linear approximation of the dispersion curve of n_{eff} (constant group index), the free spectral range is given by

$$\Delta \lambda_{FSR} = \frac{\lambda^2}{2 n_g L}, \tag{5.6}$$

still with L the half ring roundtrip.

Q factor and Finesse   The 3 dB bandwidth is given by

$$\delta \lambda_{3dB} = \frac{\left(1 - \tau^2 A^2\right) \lambda^2}{\tau A \pi 2 L n_g}. \tag{5.7}$$

From 5.7, the Q factor can be calculated as

$$Q = \frac{\tau A \pi 2 L n_g}{\left(1 - \tau^2 A^2\right) \lambda}. \tag{5.8}$$

One often distinguishes the loaded and unloaded Q, where the unloaded Q is related to the resonator would it not be coupled to waveguides. In the following, loaded Q values are always meant. The finesse $\mathcal{F}$ is the number of times the 3 dB bandwidth fits in the FSR: $\mathcal{F} = \frac{\Delta f_{FSR}}{\Delta f_{3dB}}$.

Phase response, group delay, dispersion and dynamic requirements

Both drop and pass ports experience a large phase jump around resonance, with the drop phase going through zero and the throughput phase being undefined at resonance. For a lossy resonator, the through phase goes through zero as well, and the phase change is less abrupt in both ports. The group delay reaches a maximum value at resonance for the drop port. For the through port it is negative and loses its meaning, indicating the

high pulse distortion at resonance in the remaining signal in the through port. This is illustrated in figure 5.2.

Successful filtering of communication signals puts requirements on the filter bandwidth, the roundtrip time and the group velocity dispersion. For 10Gbps NRZ signals, the 3 dB transmission bandwidth of the drop port should be around 20 GHz or 0.16 nm. Communication at 40 Gbps increases the bandwidth requirement to 80GHz or 0.64 nm. In order to avoid aliasing, the pulse width should be significantly longer than the roundtrip time in the ring, putting a lower limit to the ring circumference. At 10 Gbps and 40 Gbps, the pulse width is maximum 100 ps and 25 ps, respectively. Given a group delay τ_n normalised to the roundtrip delay, the group velocity dispersion (GVD) of the drop port is given by [64], see chapter 4:

$$D = -c \left(\frac{T}{\lambda}\right)^2 D_n = -c \left(\frac{T}{\lambda}\right)^2 \frac{d\tau_n}{d\nu} \tag{5.9}$$

with $\nu = f/FSR$ the normalised frequency. For a single pole filter with a pole magnitude r, the normalised group delay is $\tau_n = \frac{r}{1-r}$ [64]. In other words, the group delay at resonance grows with increasing pole magnitude and thus the group delay and GVD behaviour depend much on the cavity loss. Consider a constant coupling between waveguides and resonator. If the (straight waveguide) propagation loss dominates the excess loss due to bending, the cavity loss decreases linearly with decreasing radius, and the normalised group delay increases, as well as the normalised GVD D_n. However, the actual GVD is proportional to the square of the roundtrip time, counteracting the increase of D_n for smaller rings. If on the other hand the excess bending loss dominates as it is for very small rings, pole magnitude decreases with decreasing radius and both the normalised group delay and normalised and actual GVD decrease for decreasing radius. Therefore, we expect a maximum of the normalised GVD as a function of radius, but a monotonously decreasing actual GVD for decreasing radius. The expected figures for SOI ring resonators are given in the next section.

Enhancing the coupling

Racetrack resonators   With a circular resonator it is difficult to reach a high coupling coefficient if the gap size is limited by lithography. A possible solution is to add a straight coupling section, so that the resonator looks like a racetrack. Doing this adds a penalty on the roundtrip loss due to the longer cavity compared to the circular resonator with the same

radius. Coupling should therefore be slightly higher to reach the same extinction ratio at resonance. In this situation, the finesse will be slightly lower. The major drawback of this scheme is the decrease of FSR however, and the according decrease of the 3 dB bandwidth, which leads to a higher Q.

Bend-coupled resonators While the racetrack configuration enhances coupling between bus and ring, a drawback is the FSR becomes smaller. Another possibility is to use bus waveguides bent around a circular resonator [67]. In this way, the coupling can be larger without changing the FSR. For proper phase matching between the coupled waveguides, the widths need to be optimised as the bending radius for ring and bus is different. The coupling can be controlled both by the distance between bus and ring, and by the angle over which both are coupled.

Enhancing the throughput extinction With a symmetrically coupled lossy ring resonator, it is difficult to achieve a high extinction ratio. By making the coupling to the bus waveguides asymmetric however, Q factor can be traded for extinction ratio [68]. With properly set coupling coefficients, critical coupling can be reached and all power is extracted from the through port, while power transferred to the drop port is enhanced. Actually, the lower Q factor is not a drawback for communication applications. However, the transmission in the other direction, from the add to the through port, is reduced. Therefore, this solution can only be used with channel drop only filters.

5.1.2   Simulation results

In order to predict the performance with the current technology, and the achievable performance with future technology, simple simulations of the ring transmission were done, based on the frequency domain model from the previous section. For the current technology, a propagation loss of 0.4 dB/mm was assumed, corresponding to wire widths between 450 and 500 nm, and an excess loss due to bending of 4 times the measured excess loss for a 90° bend coupled to two straight wires (section 2.6.1). This assumption can be discussed, but here we consider that this loss includes in reality both the bending loss and the excess loss in the coupling section.

In order to assess possible performance enhancement, a reduced propagation loss down to 0.1 dB/mm and a 4 times smaller excess bend loss

were assumed, corresponding to goals that should be possible to meet with deep UV lithography in the near future.

Static parameters

The at-resonance transmissions, extinction ratios, finesse, Q and 3 dB bandwidth were calculated, as expected for the current technology. These parameters are given in figure 5.3 for a 5 µm radius ring as a function of κ. The Q factor is limited by the excess bend losses. However, for communication purposes, the bandwidth, finesse and extinction are more important parameters, that should be optimised by choosing a proper power coupling. About 0.4% of coupled power is needed in order to have a larger drop than through transmission at resonance, and 4% to reach an extinction ratio of 20 dB. The off-resonance extinction is then 34.5 dB, the finesse 76, the Q around 7000 and the 3 dB bandwidth is 0.22 nm or 28 GHz, which is enough to accommodate for 10 Gbps communication channels. This is shown in table 5.1, together with the parameters for a ring with 3 µm radius. However, a coupling of the order of 4% is deemed not feasible with a 200 nm wide gap in the coupler.

Figures 5.4 and 5.5 give a more complete view, plotting the same parameters as a function of radius for different coupling coefficients. For feasible coupling coefficients $\kappa = 0.10 \sim 0.15$, a Q factor higher than 10000 and a high finesse can be achieved, even for small rings down to 2 µm radius. Off-resonance, not much power is transferred to the drop port, so a high extinction ratio is easily achieved. However, the drop efficiency and extinction ratio at resonance are poor, with more power staying in the through port than transmitted to the drop port. Therefore, with the current technology, circular ring resonators are not feasible for actual communication purposes, but the high Q can be exploited for example for sensing purposes.

With better loss figures, a drastic improvement in the drop efficiency and at-resonance extinction ratio can be obtained. This is shown in 5.6. Even with a power coupling coefficient of just 2%, a 20 dB extinction would be possible with a 5 µm radius. The 3 dB bandwidth would still support 10 Gbps operation even with a single ring resonator. As the losses are thought to be mainly due to the straight to bend interface (see section 2.6.1) and the coupling region, structural optimisation and better fabrication are needed to achieve this.

Racetrack The coupling can be enhanced by using a racetrack resonator. The longer cavity experiences more loss than the circular resonator with
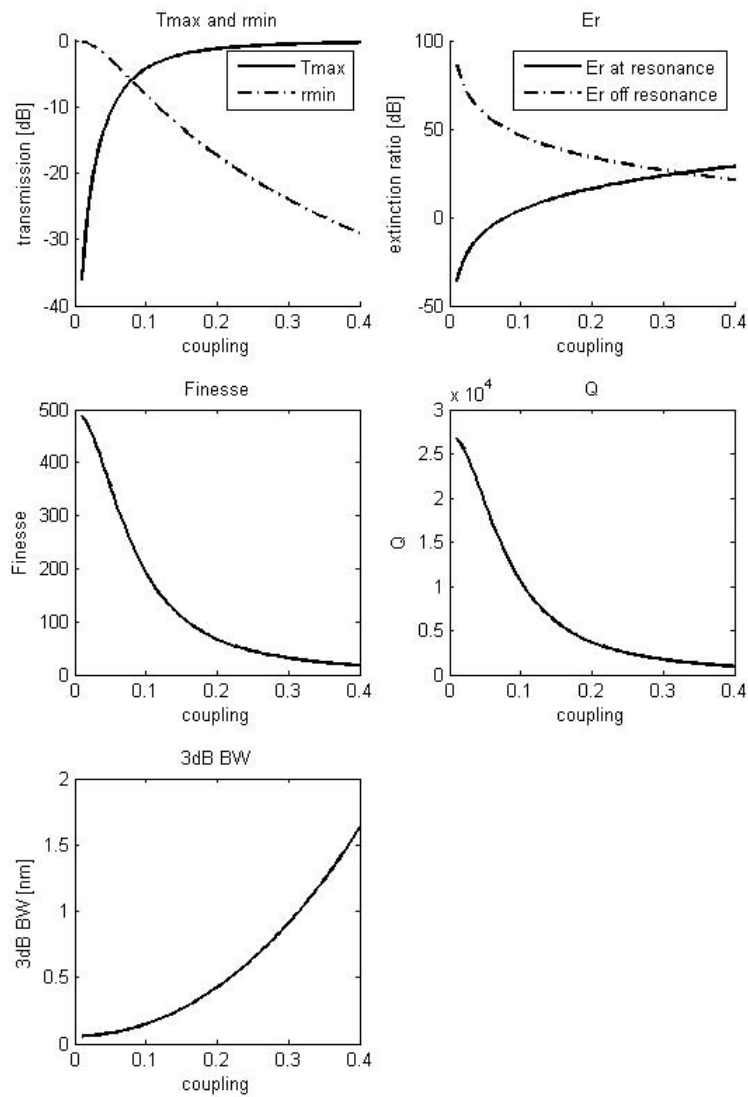
Figure 5.3: Expected performance with current technology of a ring resonator add-drop filter with 5 µm radius, as a function of amplitude coupling coefficient. Calculations based on the steady-state circuit model with realistic parameters for the current technology.

(a) Drop efficiency

(b) Extinction ratio at resonance

(c) Extinction ratio at df=100 GHz

(d) Extinction ratio off resonance
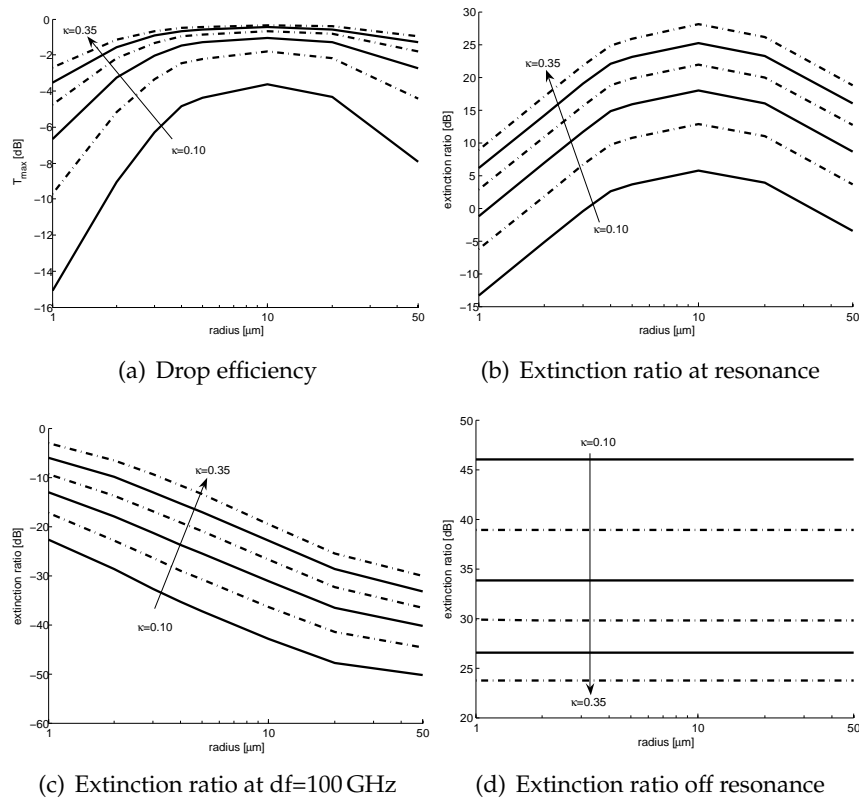
Figure 5.4: Expected performance with current technology for circular ring resonators as a function of radius, for varying coupling coefficient. Calculations based on the steady-state circuit model with realistic parameters for the current technology. Efficiency and extinction ratios in the pass port at resonance, in the drop port off resonance, and in the drop port at a 100 GHz shift from the resonance.

(a) 3 dB bandwidth

(b) Q factor

(c) Finesse
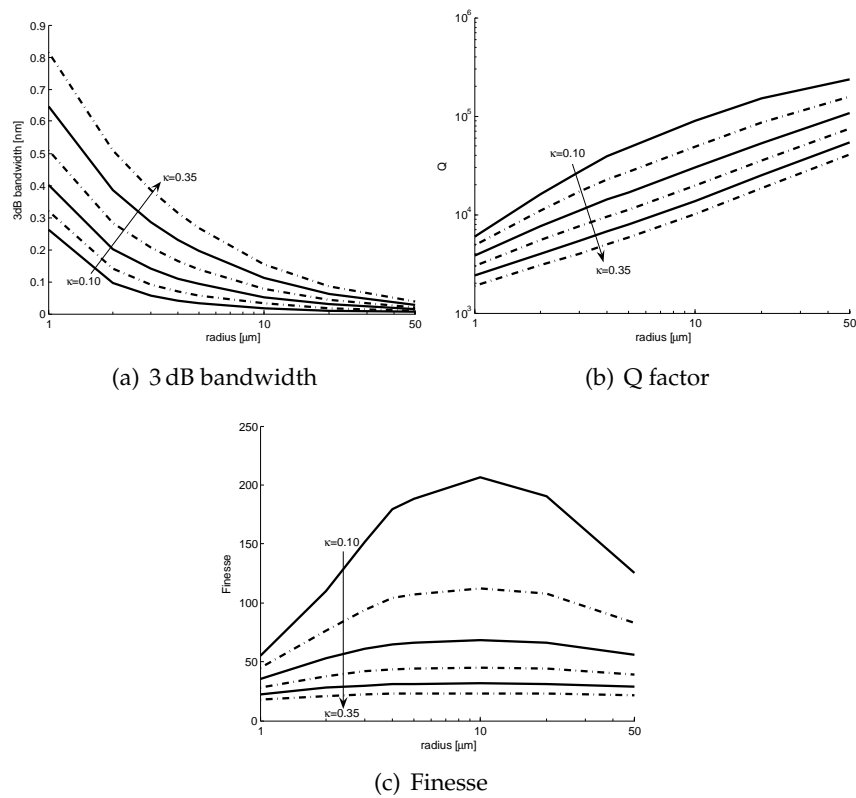
Figure 5.5: Expected performance with current technology for circular ring resonators as a function of radius, for varying coupling coefficient. Calculations based on the steady-state circuit model with realistic parameters for the current technology. Bandwidth, Q factor and finesse.

(a) Drop efficiency

(b) Extinction ratio at resonance

(c) 3 dB bandwidth

(d) Q factor

(e) Finesse
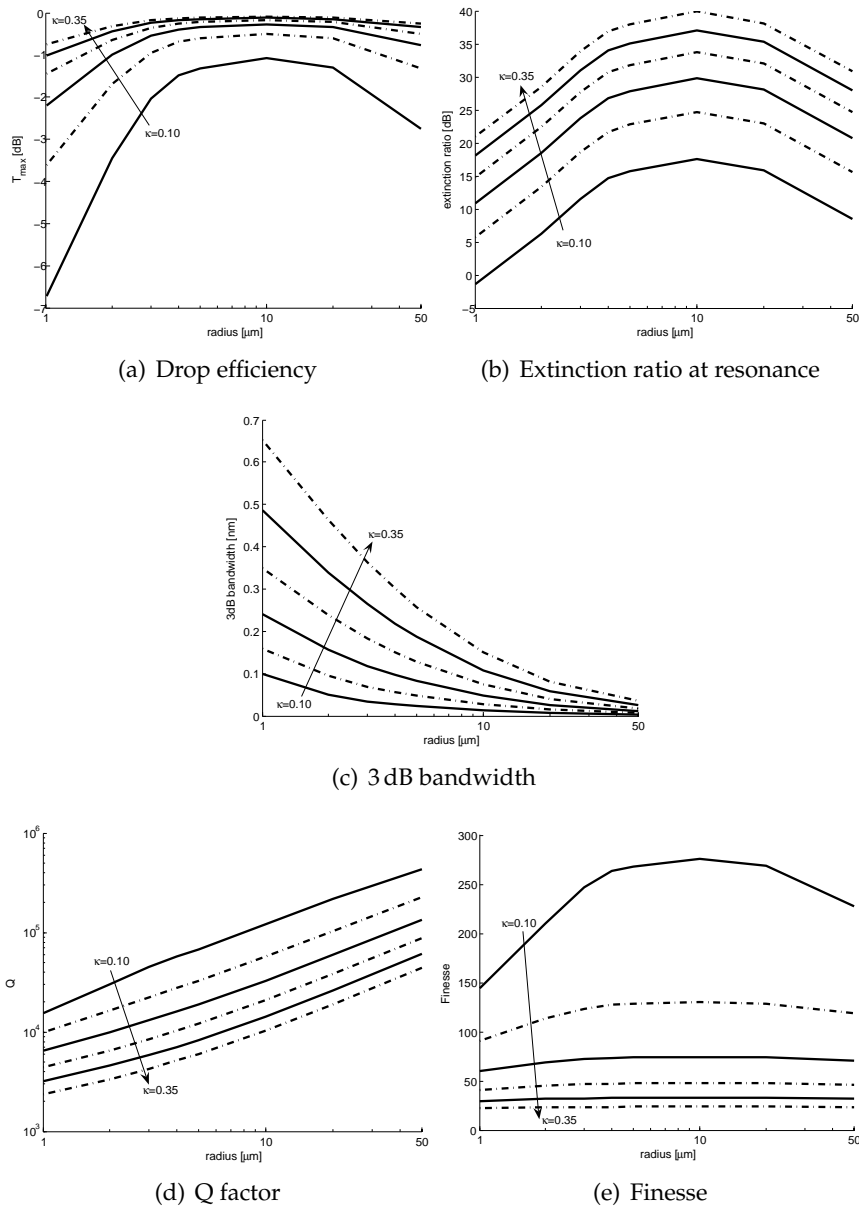
Figure 5.6: Expected performance with future technology ($\times 1/4$ bend and propagation losses) for circular ring resonators as a function of radius, for varying coupling coefficient. Calculations based on the steady-state circuit model with realistic parameters for the current technology.

R (µm)	L (µm)	roundtrip loss (dB)	κ	κ^2	E_r (dB)	E_{ro} (dB)	$\mathcal{F}$	Q	3 dB BW (nm)
5	0	0.0326	0.19	0.037	20	34.5	76	6950	0.22
3	0	0.0475	0.23	0.052	20	31.3	52.5	2870	0.54
5	6	0.0374	0.20	0.042	20	33.5	67	8470	0.18

Table 5.1: Projected performance of single ring channel drop filters with current technology, with a $E_r = 20\,dB$ defining the desired κ : circular resonator with 5 and 3 µm radius, and racetrack resonator with 5µm radius. R,L and κ are not related in this table.

the same radius, and thus a little more coupling is needed. However, with a relatively short straight section of a few µm, a power coupling of 4% to 12% should be feasible. The projected performance of such a ring is listed in table 5.1. With a similar power coupling as for a circular ring, a 20 dB extinction ratio is reached, albeit with a lower finesse, lower 3 dB bandwidth and higher Q.

Expected dynamic behaviour

For the small rings used here, with a group index of around 4.5, the round-trip time is around 100 fs to 500 fs. This is beneficial for the dispersion, as it scales with the squared roundtrip time. Also, it is not a limiting factor on the pulse width. For these roundtrip times, the GVD in the drop port is given by $D = 1.25 \cdot 10^{-3} D_n$ ps/nm to $0.031 D_n$ ps/nm. With the losses and coupling coefficients needed, the pole magnitude of the drop port transmission is 0.9 to 0.95, leading to a normalised group delay at resonance of 10 to 20. With a roundtrip time of 100 to 500 fs, the group delay at resonance will therefore be of the order of 5 to 10 ps.

Figure 5.7(a) shows the roundtrip loss for a racetrack with R=5 µm and a 6 µm coupling length, as a function of radius, calculated from the measurement results in chapter 2. The coupling was kept constant, and the extinction ratio at resonance is higher than 15 dB for R=3 µm to 30 µm. The loss has a minimum around a radius of 10 µm. With a constant κ, we can therefore expect a maximum extinction ratio around R=10 µm, as well as a maximum of the normalised GVD. However the actual GVD decreases with radius. This is shown in figure 5.7(b). The maximum GVD at R=5 µm reaches 37.33 ps/nm at an offset frequency $df = 0.044 \cdot \Delta f_{FSR} = 67\,GHz$ from the resonance. Closer to the resonance frequency, at $df = 40\,GHz$, the GVD is 32 ps/nm. This is equivalent to the dispersion through just

(a) Loss (b) GVD

Figure 5.7: Calculated loss and group velocity dispersion for constant coupling between ring and waveguides

1.9 km of standard single mode fibre. With a 100 GHz channel spacing, the GVD of the neighbouring channel in the through port is 19 ps/nm.

Lowering the straight waveguide propagation losses with a factor of 4 barely has an influence. Lowering the excess bend losses with a factor of 4 slightly increases the GVD for a constant coupling (larger extinction of 27 dB). Coupling should then not be reduced, as this will lead to a significantly higher GVD.

For the smaller ring radiuses down to 1 μm, the higher roundtrip losses can be compensated by higher coupling in order to keep the extinction ratio high. For a constant ring circumference, the GVD becomes smaller than in figure 5.7(b) for a radius below 5 μm, down to maximum 2 ps/nm for a racetrack with 2 μm radius, a 20 dB extinction ratio, a 2.6 THz (20.8 nm) FSR but necessitating a (high) 15% power coupling. Up to 40 GHz from the resonance, the GVD is only 0.6ps/nm. With a 100 GHz channel spacing, the neighbouring channel would experience a 40 ps/nm dispersion in the through port however, and the second neighbouring channel 6 ps/nm. In order to achieve a 15% power coupling with a realistic gap width, the coupling length may be made longer, with only a small penalty on the GVD and FSR decreased with a couple of nm.

As an example, a demultiplexer consisting of single ring channel drop filters on a single input bus will suffer a dispersion of less than 47 ps/nm for all channels, which is low enough for a system with a 10 Gbps datarate and a 100 GHz channel spacing without needing dispersion compensa-

tion. For a 40 Gbps system, dispersion compensation would be needed. The same holds for the example with a 5 μm radius above. However, with higher order ring filters, a similar system will always need dispersion compensation. This can be done through specially tuned ring resonators on the same chip.

Conclusion

From the previous data, it is clear a high coupling, of the order of 5% of coupled power, is needed in order to achieve an acceptable extinction ratio at resonance. As the gaps are limited to about 200 nm (see chapter 3), the coupling length therefore needs to be long enough. One possibility is to resort to a racetrack resonator, which however reduces the FSR and 3 dB bandwidth compared to a circular ring resonator with the same radius.

5.1.3   Experimental results

In initial designs (PICCO02, PICCO04, PICMOS01), the gap between bus and resonator waveguides was limited to 200 nm on the mask. Due to proximity effects, after fabrication the gap is wider (typically 250 nm) and waveguides are narrower in the coupling region. In order to obtain gaps of the order of 200 nm, the lithography illumination dose was reduced. However, both isolated and coupler waveguides become broader when doing this. Therefore, the influence of the illumination dose on the coupling coefficient is not straightforward. In later designs (PICMOS02), proximity corrections were included, with a reduced gap of 160 nm and a broader linewidth in the coupling region.

Circular resonator

Ring resonators with a radius of 3,4 and 5μm were fabricated. The waveguides were designed 450 nm wide, with a gap between bus and ring waveguides of 200 nm. Figure 5.8 shows transmission spectra of 3 and 5μm radius ring resonators. The ring with R=3μm has a wide FSR of about 29 nm, and a channel drop efficiency of -16 dB. The FSR of the ring with 5μm radius is 16.5 to 17.9 nm, with a maximum drop efficiency of about -13 dB. For both rings, the dropped power off-resonance is below the noise floor.

The coupling is clearly not high enough to compensate for the cavity losses, so drop efficiency is low. While these filters cannot be used for

(a) R=3µm   (b) R=5µm

Figure 5.8: Normalised transmission spectra of ring resonators. A noise floor is visible in the drop port due to the limited transmission of the fibre couplers far from the optimum.

actual channel dropping purposes, other applications are possible such as channel monitoring.

However, as can be seen from a finer scan on figure 5.9(a), the resonances are split. This can be explained by contradirectional coupling through the roughness of the two degenerate modes propagating in opposite directions. Little et al. [69] showed that the backreflections, accumulated in a lumped element, only need a reflectivity amplitude of the order of κ^2 in order for considerable resonance splitting to appear. The measured transmissions from figure 5.8 were fit to the model of Little et al., arriving at the fitted parameters in table 5.2. Comparing the wires here with those from [69], we see the reflectivities are very close for short correlation lengths of the roughness. The reflectivity levels from table 5.2 can be explained by a 10 nm to 30 nm correlation length and a σ_d of 1.6 nm to 2.7 nm for R=5µm and $\sigma_d = 2.4\,nm$ to $4.0\,nm$ for R=3µm. Another possibility is roughness with a correlation length of 200 nm to 250 nm with a σ_d of the order of 5 nm. Correlation lengths in the 60 to 150 nm range would generate a much larger reflectivity.

The peak splitting is only significant when the reflectivity amplitude is of the order of κ^2 or higher. By enhancing the coupling, the resonance splitting should therefore disappear. This was tested by adding a BCB[2] top cladding in order to lower the index contrast in the coupler. Resonance splitting indeed disappeared for most resonances, as shown in figure 5.9(b). Even with such a higher index top cladding though, peak splitting still appeared for some resonances.

---

[2]Benzo-cyclo-buthene, a class of spin-on polymers used in CMOS

Figure 5.9: a. Detail of figure 5.8(b) b. Same, with BCB overcladding

R	resonance	roundtrip loss	coupled power	backreflected power
3	1572.25 nm	0.0396 dB	0.24%	0.01%
5	1554.75 nm	0.0257 dB	0.18%	0.0077%

Table 5.2: Parameters of the split resonances in

Figure 5.10: Simulated power coupling as a function of gap width for a 5 µm radius ring resonator. (Simulation performed by I. Christiaens [70])

The group index of a wire rises with wavelength. However, the free spectral range is quadratic in wavelength, compensating for the larger group index. Therefore, the FSR becomes larger at longer wavelengths.

Fitting the formulas to the drop transmission from fig. 5.9(b) gives a power coupling of just 1%, and a total roundtrip loss of 0.07dB ($A^4 = 0.98$). Thus, roundtrip losses seem higher than expected from the loss measurements in chapter 2, in the case of this R=5 μm ring 0.0225 dB, or 0.72 dB/mm higher. This could partly be explained by narrower waveguides, but top near-field measurements in section 5.3 point in the direction of an excess loss around the coupling region, which would lead to an excess loss of 0.0112 dB per coupler in this case.

Another possible source of lower than expected drop efficiency is a deviation of the coupling region from the ideal behaviour. Due to the fast transition from an isolated bend mode to a coupled mode in the coupler, the coupled supermodes are not exited in the properly balanced way, leading to a phase difference between the outputs that differs from $\pi/2$. This not only influences the resonance wavelength, but also distorts the power transmission around resonance. However, due to the improper phase matching that happens in a circular resonator coupled to straight waveguides anyway, this effect is difficult to model. In a racetrack resonator, the nonunitary transfer is easier to model and can be shown to have a large effect in certain cases [71]. The same effect could be explained by the supermodes not having an equal attenuation in the coupling region, as their overlap with sidewall roughness or other imperfections is different.

Racetrack

As long as losses are too high, coupling needs to be enhanced in order to arrive at a high drop efficiency. This can be done through a racetrack configuration. Q factor and FSR will be lower, though. Figure 5.11 shows transmission spectra for racetracks with 1 μm, 2 μm and 5 μm radius. The racetrack with R=1μm still has a Q higher than 2000, and a 44 nm FSR. The high losses of 0.18 dB per roundtrip (fitted), and still limited coupling (1% fitted), still limit the drop efficiency. Still, it shows bends can be made with very low loss. With a 2 μm radius, the drop efficiency is already a lot better, but still E_r is negative. The Q is higher than 6000, coupling is fitted to be around 1.2%, the total roundtrip loss is 0.06 dB. With 5 μm radius and a longer coupling length, coupling is high enough to compensate for losses and an extinction ratio better than 20 dB is obtained. The Q is still higher than 3000, the finesse is 26 within a 11.5 nm FSR. The fit delivers a

9% power coupling, and a 0.12 dB of roundtrip loss. The waveguides in Fig. 5.11(c) have smaller widths than those in Figures 5.11(a) and 5.11(b), and the cavity is longer, explaining the higher losses of the R=5µm racetrack compared to the R=2µm racetrack.

The first component mask (PICCO02) contained ring resonators with more narrow waveguides (400 nm designed) and 200 nm designed gaps. These waveguides experience a higher loss but coupling is higher. A ring resonator with 5 µm radius and 3 µm coupling length had a Q of 3000 and a FSR of 14 nm. The measured spectra are shown in 5.12. An SEM picture of the resonator is shown in figure 5.13.

The fifth component mask (PICMOS02) contained racetracks with an equal radius of 4 µm, but increasing coupling lengths of 3, 5, 7 and 8 µm. The fitted coupling increased with increasing coupling length from about 1.9% to 9.1%, while the Q decreased from 10500 to 3200. At the same time, the extinction ratio increased from 7.3 dB to 14.23 dB for 3 to 7 µm coupling length, but decreased again for L=8 µm, which can be explained by an increasing roundtrip loss, from 0.035 dB to 0.106 dB (fitted).

One detrimental effect for ring and racetrack resonators is distortion due to non-ideal coupler transfer. Due to loss in the coupler or asymmetric excitation of the supermodes due to the fast change of mode profiles just near the coupler, the phase relation between the input wave and the resonator mode is distorted. Constructive interference in the pass port happens at frequencies slightly differing from the resonance frequencies. In a lumped circuit model, this can be modelled as the output ports of the coupler not being in quadrature as for an ideal coupler. The effect is not only a shift, but also an asymmetry in the spectrum. This is illustrated by the measured spectra in figure 5.14. With the loss levels mentioned previously, the steady-state lumped model shows a deviation of $\frac{\pi}{20}$ to $\frac{\pi}{30}$ is enough to explain the small but nevertheless significant measured distortion. However, the lumped model is only a very rough approximation in the case of small, high index-contrast resonators. As the cause of the distortion is twofold, this problem can be tackled on two fronts. Processing technology can be improved in order to reduce the losses, and the transition between the bend and the coupled waveguide region can be made smoother. The latter of course has an adverse effect on the flexibility of the ring structure and the specs that can be reached. A important note is that better losses should come hand in hand with better control over the coupling, as higher Q structures are also more sensitive.

A second problem is the significant variation of the coupling coefficient with wavelength. This can be seen in figures 5.11 and 5.12, where

(a) R=1µm,L=3.14µm

(b) R=2µm,L=3.14µm

(c) R=5µm,L=6.28µm

Figure 5.11: Transmission spectra of racetrack resonator channel drop filters

Figure 5.12: Measured transmission spectrum of a racetrack resonator with R=5µm, L=3µm, with narrower waveguides (400 nm)

Figure 5.13: SEM picture of a racetrack resonator with 5 µm radius and a 3 µm straight coupling section.

(a) $R=2\,\mu$, $L_{coup} = 3.14\,\mu$m   (b) $R=5\,\mu$, $L_{coup} = 6.28\,\mu$m

Figure 5.14: Two measured examples of relative shift and asymmetry of the drop and pass port transmission, most probably due to non-ideal coupler transfer, leading to a distorted phase relation between the resonator resonance and the exciting wave.

the roll-off of the resonances becomes worse with wavelength. The wavelength dependency is intrinsic to the coupler structure in SOI. Therefore, ring resonator design that is generic with respect to wavelength is difficult, definitely if the losses are reduced in the future.

Example of a phase transfer and group delay behaviour

For a racetrack resonator with a 5 μm radius and 6.28 μm coupling length, the phase transfer and the group delay were calculated from the dropped power transmission spectrum. This can be easily done be realizing that the phase and magnitude are related by a Hilbert transform (as holds for every causal system with a real impulse response) [64]. For a minimum-phase, stable filter, such as the drop transfer of the passive resonators studied here, it is possible to calculate the phase response from the magnitude transmission. The result is shown in figure 5.15. The maximum group delay is around 13 ps, corresponding well with the ring roundtrip 660 fs and a pole magnitude of 0.95 (normalised group delay of 19) fitted from the power transmission spectrum.

Bend-coupled circular resonator

Figure 5.16 shows an SEM picture and the transmission spectrum of a bend-coupled ring resonator with 5 μm radius. The Q is higher than 5000,

Figure 5.15: Measured normalised power transfer and calculated phase transfer and group delay of the drop port of a racetrack resonator with R=5 μm, L_{coup}=6.28 μm
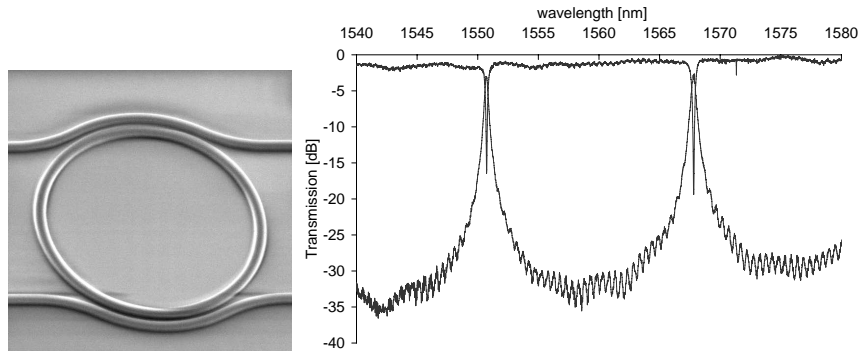
Figure 5.16: SEM picture and transmission spectrum of a ring resonator coupled to bent waveguides. R=5 μm

the FSR is 17.6 nm ($\mathcal{F}$=57), the roundtrip loss 0.07dB and the power coupling about 4%. The measured drop efficiency is limited to about -2 dB.

Operation in other frequency bands

With gaps of the order of 200 nm and waveguide widths around 450 nm, coupling and transmission can be made relatively good around 1550 nm wavelength for the TE_{00} mode. At longer wavelengths, the TM_{00} mode is less confined, coupling becomes higher, Q degrades and the extinction ratio at resonance increases. Therefore, it might be beneficial to operate at longer wavelengths. This is illustrated in figure 5.17(a), which was obtained with facet coupling through AR coated facets.

At shorter wavelengths, the TE_{00} mode becomes very well confined and coupling decreases, up to the point where the resonance is fully depleted. However, the TM_{00} mode becomes reasonably well confined, experiences less substrate leakage and can start resonating. This is shown by the measurement in figure 5.17(b), which was performed at the Valencia Nanophotonics Technology centre with a system of lasers tunable from 1260 to 1600 nm. The facets were not coated and the setup had some alignment problems at the time, explaining the noisy spectrum and the poor transmission. Nevertheless, resonances can be clearly distinguished in the 1260-1350 nm range indicating TM operation.

(a) racetrack with R=5 μm, around 1600 nm   (b) racetrack with R=4 μm

Figure 5.17: Transmission spectra of racetracks at 1600 nm and 1260-1500 nm.

5.2 Single-ring all-pass filters

5.2.1 Steady-state model

With the same notations as for a channel drop filter, the equations for the transmission of an all-pass filter are

$$T_p = \frac{T_{min} + F \cdot \sin\frac{\phi}{2}^2}{1 + F \cdot \sin\frac{\phi}{2}^2} \tag{5.10}$$

$$T_{min} = \frac{\left(\tau - A^2\right)^2}{\left(1 - \tau A^2\right)^2} \tag{5.11}$$

$$F = \frac{4\tau A^2}{\left(1 - \tau A^2\right)^2} \tag{5.12}$$

$$\Phi = \arctan\frac{A\sin\phi}{\tau - A\cos\phi} \tag{5.13}$$

with Φ the phase response.

A lossless ring (A=1) has a unity transmission but a phase response as depicted in figure 5.18(a). For a lossy cavity, the input signal is fully extincted if the ring is critically coupled: the coupled power is equal to the power lost in the ring: $\tau^2 = A^4$ or $\kappa^2 = \left(1 - A^2\right)^2$. For lower coupling (undercoupled ring) and higher coupling (overcoupled ring), the extinction ratio is finite.

(a) lossless all-pass filter   (b) lossy all-pass filter

Figure 5.18: Power and phase transmission spectra of single ringresonator all-pass filter with and without loss.

5.2.2 Projected behaviour

All-pass filters can be used for their phase transfer properties or their at resonance extinction when loss is introduced in the resonator. In this work, all-pass filters were fabricated mainly to extract the cavity losses. With a roundtrip loss of 0.0374dB (A=0.9978), a power coupling of 0.88% is needed for critical coupling. Due to the high Q, the extinction ratio is however very sensitive to small coupling variations. A marginally higher power coupling of 0.96% leads to a -25 dB extinction, a coupling of 1.05% to -20 dB.

5.2.3 Experimental results

Figure 5.19(a) shows the transmission of a racetrack with a 5 µm radius, a 4 µm coupling length, and a gap of 180 nm (designed) between racetrack and waveguide. As the extinction is nearly 30 dB, the resonator operates near critical coupling. The fitted roundtrip loss is around 0.15 dB, again much higher than expected, but the power coupling reaches about 4%.

Fitting of losses and coupling is however difficult due to the asymmetric transmission spectrum, which is clearly visible. This asymmetry can have two causes.

One possible cause has already been seen in racetrack resonators in section 5.1.3: distortion of the phase transfer of the coupler. Again, here a phase deviation of the order of $\frac{\pi}{30}$ to $\frac{\pi}{20}$ would be needed to explain

(a) P$_{in}$=2 mW   (b) Varying input power
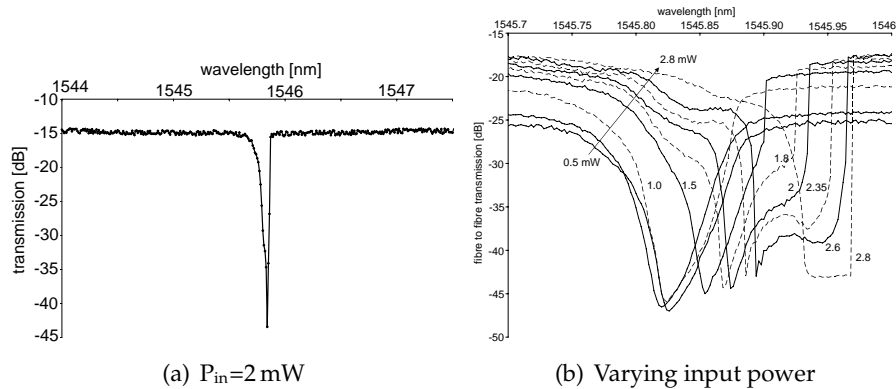
Figure 5.19: Fibre to fibre transmission spectrum of a racetrack coupled a single waveguide. R=5 μm, coupling length 4 μm.

the order of magnitude of the asymmetry in the case of the ring from figure 5.19. However, other effects are dominant here: secondary thermal effects originating from non-linear processes (two-photon absorption, and the resulting (linear) free carrier absorption and surface state absorption). The same spectrum is shown for various input powers in figure 5.19(b). Already at modest power levels in the cavity, a thermal effect may create a typical asymmetric transmission spectrum, and at higher power levels bistable behaviour is obtained [24]. Near critical coupling, the field enhancement of the resonator is at its highest. The fitted loss and coupling for the resonator from figure 5.19 give a field enhancement of 5.3. With 1 mW optical power in the bus waveguide, this gives power in the ring of the order of 25mW. At this level, the thermal effect may indeed be non-negligible. For increasing power, the spectrum is increasingly distorted. This thus happens already at input power levels as low as 2 mW.

5.3   SNOM measurements

Scanning near-field optical microscopy (SNOM) can reveal a lot of information not directly apparent from transmission measurements, by imaging the near field above a structure. As the near field contains information about high spatial frequencies that is lost in the far field, the details of the optical field just above the structure can be imaged [72]. A fine fibre tip is scanned over the component, detecting the near field, while light is launched into the component through facet coupling. The near field is a

combination of the evanescent tails of guided modes and scattered field that will become the propagating far field. Therefore, care is needed while interpreting the figures. Figure 5.20 shows SNOM images of some race-track resonators. These measurements were performed by colleagues at INSA Lyon. Resolution of the images between 30 nm and 36 nm. Distance between the fibre tip and the sample is 10 nm, kept constant by a control system.

In figure 5.20(a), the bottom right waveguide is the input, bottom left is the dropped output. At-resonance, we can see the field enhancement in the ring. Scattering is apparent at the bottom side of the ring, as the first (top) half of the ring would carry a higher intensity. In figure 5.20(b), a racetrack with 1 µm radius is shown. Here, we can clearly see the enhanced scattering in the transition region between coupling region and bend. This scattering can have three possible origins. Due to optical proximity effects, the isolated wire is broader than the double wire, with a short inadiabatic transition between them. Secondly, the 1 µm bend forms far from a slow transition from coupled wires to isolated bend, creating mode mismatch. Thirdly, accidentally there may be extra scattering defects in that transition region. From the detail in figure 5.20(c), we can see a standing wave pattern resulting from the contradirectional coupling, and scattering defects at the top of the sidewalls of the wires, which is consistent with the presence of microholes as shown in section 3.2.3.

5.4 Ring-resonator-based demultiplexer

By mounting ring resonators with differing resonance wavelengths on a single input bus waveguide, one obtains a demultiplexer. A possible advantage of such a ring-resonator based components is that the wavelength channels can be set at will, without a need for constant channel spacings. Disadvantages are the difficulty of control over the resonance wavelengths, and the need for more complex series coupled resonators to achieve a large enough bandwidth for actual WDM applications. Figure 5.21 shows the transmission spectrum of a demultiplexer using a different ring radius for each wavelength (6, 6.02, 6.04, 6.08µm radius). The resonances are in the expected order, but as we will see in chapter 8 this only happens for some components.

(a) R=5 µm, L=6.28 µm. left: at resonance, right: off resonance

(b) R=1 µm, L=3.14 µm. left: topographical image, right: SNOM

(c) detail of coupling section

Figure 5.20: SNOM measurements of racetrack resonators and topological images corresponding to the characterized structures. The measured field intensity is plotted with a resolution of about 30 nm. Measurements performed at INSA Lyon.

(a) Fibre-to-fibre transmission spectrum

(b) SEM picture

Figure 5.21: Ring resonator demultiplexer. Each ring has a slightly different radius

5.5   Higher order filters

Although performance of the single ring filter is critical, its use is limited for actual WDM. The bandwidth and roll-off can be improved by using higher order filters. A few ring-resonator based higher order filters were designed: second order and third order lattice filter, and a third order Butterworth (ARMA, see chapter 4) filter.

5.5.1   Ring lattice filters

For filters based on ring resonators only, mainly two options are possible. Either one uses a cascade of uncoupled rings, in which each ring can easily be tailored with respect to coupling and cavity length, or a lattice of coupled rings. In contrast to the cascade filter, the lattice filter can be designed to have a 100% transmission in the passband for a lossless system. The cascade filter can either be optimised for high transmission and small bandwidth or the other way round. However, design of the lattice filter is more complex as the location of the poles are a complex function of the coupling coefficients.

   The main problem of the high index contrast is the difficulty of aligning the uncoupled ring resonances due to small waveguide variations. Therefore, fabricated structures were limited to two-stage and three-stage filters. Figure 5.22 shows the measured transmission for a second-order lattice with designed coupling coefficients of κ =0.6/0.3/0.6. The rings were racetracks with 5 µm radius and a 19.7 µm coupling length in order to achieve $\kappa = 0.6$. The gaps were designed 180 nm/210 nm/180 nm. The drop port transmission is broadband, but the loss is already 3 dB. The through port reaches a 15 dB extinction, but only in a narrow band corresponding to a single resonance. Deviations from the optimal curve are mainly due to the uncoupled resonances not being properly aligned. On other samples, this misalignment was worse.

   All fabricated third-order filters contained misaligned resonances that destroyed the transmission.

5.5.2   Single-stage ARMA filter

One of the most simple ARMA-type filters, is a Mach-Zehnder with an all-pass ring resonator in one arm and a delay in the second arm, as depicted in figure 5.23(a). This is designed as a third order Butterworth filter [64]. In other words it has maximally flat passbands and stopbands, and compared to FIR filters (see chapters 4 and 6), a much steeper roll-off for a

(a) Schematic   (b) Measured transmission
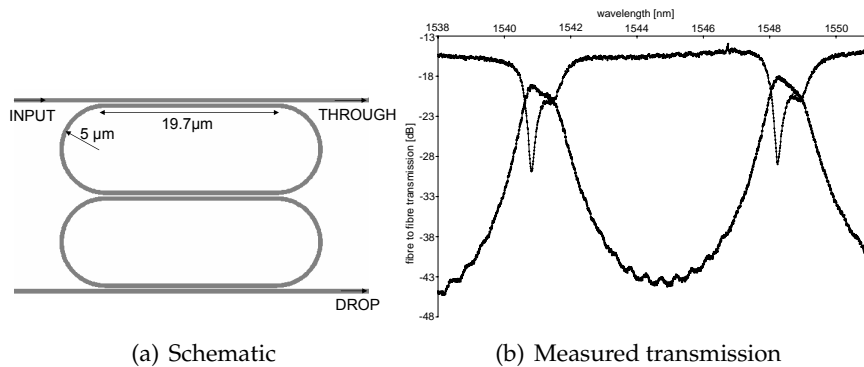
Figure 5.22: Second order ring lattice filter with two racetracks with 5 μm radius, and coupling coefficients κ =0.6,0.3 and 0.6.

smaller filter order. This filter can be used as a channel interleaver or a duplexer.

At the anti-resonant frequency, the phase transfer of the ring changes slowly. Near the resonance, the phase change is stronger. The anti-resonant frequency is aligned with the passband centre by adding a phase factor π. A flat passband can be obtained by having a constant phase difference between both arms of the MZI around the passband centre. This can be done by setting the group delay of the ring at the passband centre frequency equal to the group delay of the delay line in the second arm. One solution is to choose the phase delay in the ring to be double the group delay in the delay line [73, 74], obtaining a third order filter. This corresponds to a pole magnitude of 1/3. The power coupling needed is 89% [73]. At this small pole magnitude, the phase change near resonance is still relatively slow. If coupling differs from this critical coupling, the extinction in the stopband decreases and transmission in the passband decreases.

With SOI photonic wires, a relatively large FSR for such a filter can be achieved, enabling other applications than the very small FSR devices developed in low index contrast technology [73]. An example of this filter was designed with a FSR of 12 nm ($\approx$1500 GHz). The limiting factor to the FSR is the coupling length needed to achieve a 89% power coupling. A racetrack resonator with a 5 μm radius was used and a straight coupling length of 31 μm. The measured cross and bar transmission for the best lithography dose are shown in figure 5.23(c). Wide and relatively flat passbands are achieved, with a non-uniformity of 0.5 to 1.5 dB, and a very steep roll-off. The extinction in the cross transmission stopband is accept-

(a) Filter layout
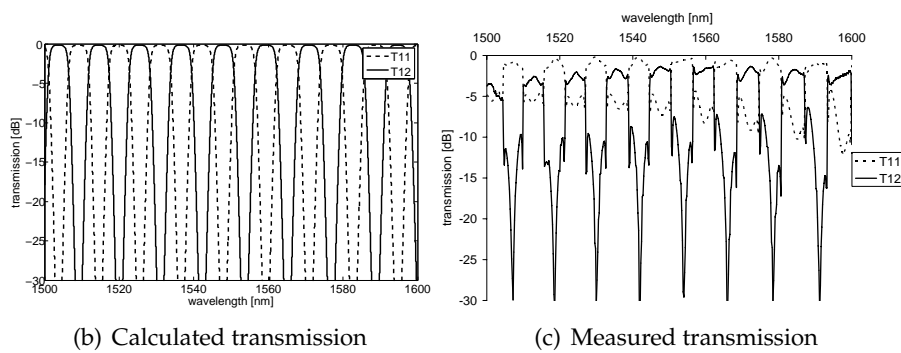
(b) Calculated transmission

(c) Measured transmission

Figure 5.23: Third order Butterworth filter based on a ring resonator in a Mach-Zehnder interferometer. Schematic, calculation with circuit model and realistic loss parameters, and measurement result.

able, at -6.8 dB near the band edges, and better than 25 dB at the stopband centre. However, in the passband, the extinction is only of the order of 2.2 to 5 dB at the edges and 6.5 dB at the passband centre. Further optimisation of the structure and phase control using heaters are needed to improve this performance.

5.6 Conclusion

Single-ring-resonator channel drop filters with acceptable extinction ratios were demonstrated in the form of racetrack and bend-coupled circular resonators. While the cavity roundtrip losses are still high, this could be compensated for by higher coupling. Circular resonators coupled to straight waveguides showed Q factors higher than 10000.

In the near future, better processing technology may lead to single ring filters that are actually suited for 10 Gbps communication channels. A smaller sidewall roughness is needed to reach this. The main problem is not the scattering loss but the backscattering to the counterpropagating mode. We have experimentally shown this problem can be alleviated by a higher coupling, but this reduces the design options. Also, the backscattering problem can not be dealt with by addition of tuning elements, in contrast to a certain degree of irreproducibility that can be compensated for by temperature or other tuning. A smaller roughness needs to come hand in hand with better control over the coupling between bus and ring waveguides, so process development needs to work on both fronts simultaneously. However, lower losses relax the minimum gap width needed and therefore enable better control over the gap.

A single-ring based demultiplexer was fabricated, and showed clearly distinguishable dropped channel wavelengths. However, control over the resonances should be improved. Two examples of higher-order filters based on ring resonators were shown: a second-order ring lattice filter with a large passband width but still rather weak extinction, and a third-order Butterworth filter with relatively flat passbands, but again with an extinction needing improvement.

These results show that if dimensional control can be improved, ring-resonator based channel drop filter devices with a footprint smaller than $0.01 \, \text{mm}^2$ can be obtained.

Chapter 6

Mach-Zehnder-based filters

In this chapter, non-resonant Mach-Zehnder based channel drop filters will be studied. Such structures, with feed-forward paths only, have large passband widths but need multiple stages in order to achieve a high stop band extinction ratio. This is in contrast to the ring resonator filters from chapter 5, which intrinsically have a large stopband extinction, but need multiple resonators to achieve a large bandwidth needed for high data rate transmission and tolerance to wavelength deviations.

Also, the high dispersion of transmission through resonant structures around the resonance wavelength necessitates dispersion compensation, whereas some types of Mach-Zehnder based filters as studied in this chapter can be linear-phase by design.

Fabrication variations have a larger impact on the stopband transmission than on the passband, generating sidelobes in the transmission spectrum which lead to crosstalk.

We have studied how such filters perform with the current generation of CMOS-based wafer-scale fabrication technology. First, the simple two-port and four-port Mach-Zehnder interferometers are covered as basic building blocks and then the design and characterisation results of lattice filters and Mach-Zehnder cascade filters are described. We have analysed the measurement results and simulated the devices on a circuit level with realistic parameters, giving insight into the relative contributions to the crosstalk and how to reduce them in the future.

6.1   Mach-Zehnder interferometer

The Mach-Zehnder interferometer (MZI) is a basic building block and can be employed as a modulator, switch and wavelength filter suited for channel interleaving.

6.1.1   Two-port MZI

Projected performance

For a MZI with single inputs and outputs (1-by-1) and a length difference between both arms, the extinction ratio of the throughput depends on the splitting ratio of the splitter and combiner, and on the relative loss between the arms. Figure 6.1(a) shows the calculated extinction ratio of a 1-by-1 MZI as a function of delay length for several loss figures. With the current propagation losses, the extinction ratio of a MZI with a 1 mm delay length is better than 25 dB, and up to 45 dB with 500 nm wide waveguides. MZI devices with a ten times larger FSR would have a loss-limited extinction of at least 10 dB better depending on the waveguide width. However, from figure 6.1(b), plotting the extinction as a function of splitter asymmetry for a 1 mm delay length with a typical current loss figure, we can see that the extinction ratio is determined by the symmetry of the splitters at a given loss level, rather than by the losses themselves. For this example, perfectly balanced splitters would already reduce the extinction to 33 dB. Coupling 0.2 dB less power to and from the more lossy arm compared to the other arm reduces the extinction to 27 dB, while an asymmetry in the opposite direction enhances extinction. For shorter delay lengths (larger FSR), the problem is smaller but still a 0.2 dB asymmetry in either direction can easily reduce the extinction with 10 dB.

A reduced propagation loss of 1 dB/cm and, more importantly a bending loss smaller than 0.01 dB, which should be achievable in the near future, would already deliver an extinction performance that is an order of magnitude better with respect to splitter asymmetry.

Measurements

MZI devices were fabricated with delay lengths from 20 μm up to 1 mm in one arm, corresponding to 27 nm down to 550 pm free spectral range respectively. Figure 6.2 shows some measured transmission spectra. The insertion loss of around 4dB is mostly due to the unoptimised Y-splitters used. An extinction ratio larger than 25 dB and reaching 30 dB is achieved with a 1 mm delay length. With the measured losses from chapter 2, the imbalance can be estimated to be smaller than 0.13 dB for both splitters. The similar maximum extinction ratio for shorter delay lengths shows the importance of loss with respect to splitter balance.

(a) as a function of delay length

(b) equicontours as a function of asymmetry of both splitters, for 1 mm delay length and 0.4 dB/mm loss
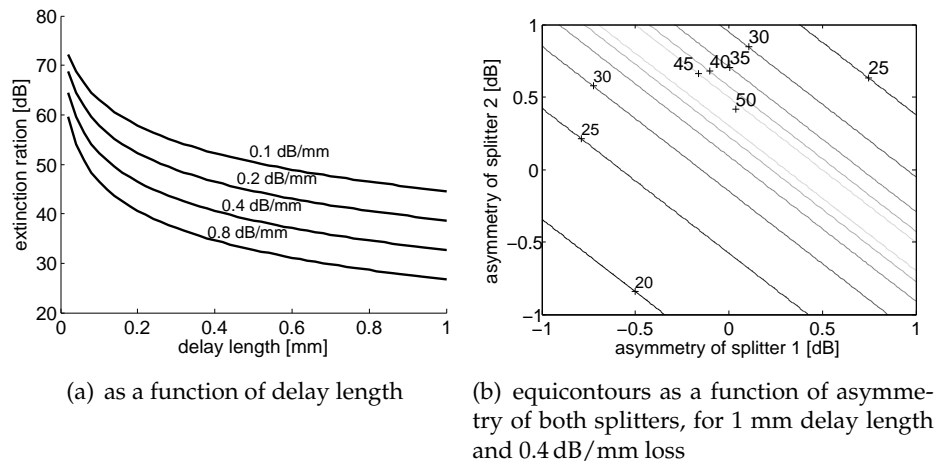
Figure 6.1: Calculated extinction ratio of 1×1 MZI devices

(a) 50 µm delay length

(b) 1 mm delay length

Figure 6.2: Measured transmission spectra of 1×1 MZI devices

(a) 50 μm delay length

(b) 200 μm delay length
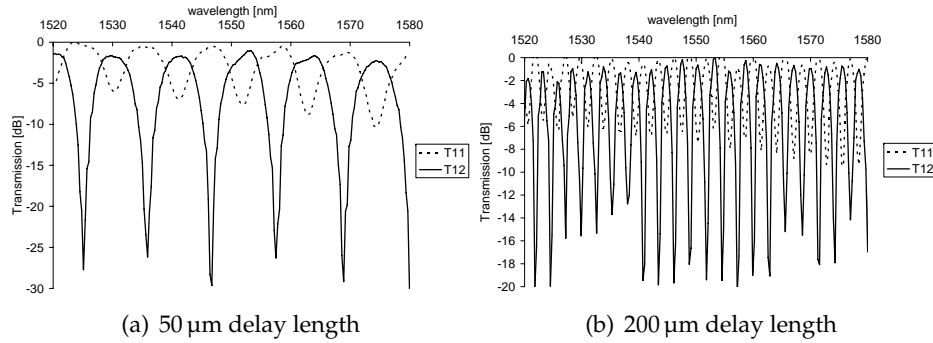
Figure 6.3: Measured transmission spectra of 2×2 MZI devices

6.1.2 Four-port MZI

For channel interleaving purposes or as a basic building block for cascaded or lattice MZI filters, a device with two inputs and two outputs is needed. Such devices were designed and fabricated with directional couplers as 3 dB splitters. Controlling the coupling of directional couplers proved not easy however, and their wavelength dependency makes the extinction between ports wavelength-dependent. Figure 6.3 shows measurement results for 2×2 MZI devices with 50 μm and 200 μm delay length. The extinction the cross transmission reaches 25 dB and 18 dB for both devices, respectively. However, in the bar transmission, extinction is only 7 dB. This indicates a serious unbalance of the coupled powers between both arms. Calculation shows that the coupled powers should be of the order of 70%/30% to explain this measurement. The measurement shown is for the best lithography dose, indicating control over the directional couplers is not good yet. In the next section, we will see this has an influence on the transmission of higher-order filters based on Mach-Zehnders.

6.1.3 Conclusion

For better extinction ratios, which can be needed for certain higher order filters based on these MZI devices, the loss should be further reduced, not because of the loss as such but because only then splitters can be perfectly balanced and the MZI is more tolerant to splitter asymmetry. Additionally, four-port power splitters with drastically better coupling balance are needed for four-port MZI devices.

Figure 6.4: SEM picture of an 11-stage lattice filter

6.2 Lattice filters

In order to enhance the extinction and roll-off, multiple MZI structures can be combined. In a lattice filter, also called cascaded Mach-Zehnder or resonant coupler, the Mach-Zehnder stages are concatenated as depicted in figure 6.4. First, the basic operation of the lattice filter will be reviewed, its advantages and the design method for obtaining a channel drop filter. Then, it is briefly discussed how to reduce phase errors in the delay lines, and measurement results of devices with 100 to 400 GHz channel spacings are shown. A discussion of the measurement results follows, intermixed with simulation results, in order to get a feeling for the order of magnitude of the fabrication deviations generating crosstalk. Based on these results, we finally try to predict the performance of SOI lattice filters in the future.

6.2.1 Basic operation and equations

Each delay is an integer multiple of a unit delay length, and the coupling coefficients of the couplers can vary in order to approximate the desired transmission. With N stages (delay lengths), N+1 light paths are possible between an input and an output port, so a filter of order N needs N delay lengths and N+1 couplers. Such a filter can be designed to drop a single or even multiple channels while the other channels go to the power complementary output.

The lattice filter can easily be described using a matrix formalism. Each lossless coupler is described by a matrix

$$S_c^i = \begin{bmatrix} \tau_i & -j\kappa_i \\ -j\kappa_i & \tau_i \end{bmatrix},$$

(6.1)

with κ the amplitude coupling and $\tau = \sqrt{1 - \kappa^2}$. The $\frac{\pi}{2}$ phase shift expressed by j only holds for an ideal lossless coupler with at least one sym-

metry plane. Moreover, the reference plane in each port of the coupling structure should be chosen so that the waveguides are uncoupled at the reference planes.

A delay section is given by the matrix

$$S_d^i = \begin{bmatrix} \mathrm{e}^{-j\beta\Delta L} & 0 \\ 0 & 1 \end{bmatrix} \tag{6.2}$$

with β the propagation constant in the wires, and ΔL the physical (unit) delay length.

The full transmission of a N^{th} order lattice filter is then given by

$$S = S_c^{N+1} \prod_{i=1}^{N} S_d^i S_c^i. \tag{6.3}$$

If the phase factors are generalised to the full complex plane, S_{11} and S_{12} can be written as polynomials in $z = |z|\,\mathrm{e}^{-j\phi_z}$. In other words, due to the discrete delay lengths used, a Z transform description is possible. The optical transmission is obtained by evaluating the polynomials on the unit circle: $z = \mathrm{e}^{j\beta\Delta L}$ (see chapter 4).

6.2.2 Advantages of lattice filters

An advantage of lattice filters is the low insertion loss, as in theory these filters can be lossless [65]. While multiport structures such as the AWG may offer a more attractive solution for routing and massive demultiplexing operations, they are often intrinsically lossy. A non-lattice form cascade of Mach-Zehnders can also be used for channel dropping purposes, but will experience more loss than a lattice filter. Also, with addition of phase shifters in each stage, arbitrary transfer functions can be realized both in amplitude and in phase, for use in wavelength filtering, dispersion equalisation and other applications. As each stage has a multiple of the unit delay, the impulse response can be represented by a finite discrete-time Fourier transform. Lattice filters bear a lot of resemblance to digital filters with a finite impulse response (FIR), composed of feed-forward paths only. It is therefore no surprise that a Z transform description of the filters is possible and digital filter design methods can be used [64]. Here, only lattice filters without phase shifters (real coefficients) are considered for 1 out of N channel selecting. An advantage of filters with real coefficients is that they are linear-phase if the coefficients are symmetrical, so they have a constant group delay and no distortion of the signal in the passband. A last advantage is the relatively easy trade-off of adding

more stages: better spectral performance (steeper roll-off, better extinction, more wavelength channels per free spectral range) comes at a larger device size, a more tight fabrication control and longer impulse and step responses[1], but this can be controlled on a per-stage basis.

6.2.3 Design

A number of design options are possible for this type of component. Using an optimisation algorithm for the coupling values, the desired transfer spectrum can be approached, starting from high-level specifications such as the tolerated crosstalk and passband ripple and the desired bandwidth. One can have a discrete set of coupling values (for instance if the couplers are implemented as MMI structures), or allow any value. A more attractive approach is to use digital filter design methods, which are well established, and then convert the digital filter coefficients to the optical coupling factors needed. The basic design steps and the relation between optical and digital filters were laid out by Jinguji and Kawachi [65]. A very good introduction to filter design in general can be found in [64].

The first step is to find a polynomial in the Z plane of a given order using a digital filter approximation method, optimising the locations of the zeros in the Z transform plane. Then the boundary conditions needed for valid optical filters are applied: power normalisation and power conservation (complementary outputs). The next step is to deduce the coupling coefficients from the polynomial coefficients. Various digital filter design methods can be used, amongst which one of the most widely used is the Parks-McClellan algorithm [66]. This minimax algorithm minimises the maximum weighted error (Chebyshev error) between the desired filter characteristic and a polynomial of a given order, leading to an equiripple in the frequency response and a steep roll-off. Numerous other approaches exist, such as several windowed methods and the least-squares approximation, which does not exhibit an equiripple, a less steep rolloff, and a smaller error over most of the passband but a larger error near the passband edges. Functions for designing digital filters with these methods are available in Matlab's Signal processing toolbox and in Scipy's signal package. An optical bandpass filter, with its centre frequency much larger than the FSR, can be designed as a digital lowpass filter.

Equiripple channel drop filters with a limited number of stages were designed with the Jinguji method [65], which was implemented in Matlab

---

[1]The impulse and step responses of a FIR filter are as long as its order, corresponding to the different total delays a signal may experience

and Python scripts. The input parameters are the number of channels, the channel spacing, the passband end (or half the required passband width), the tolerated ripple in the passband and the tolerated crosstalk (ripple in the stopband). First, the Parks-McClellan algorithm is used to obtain the polynomial $F(z)$ of one of the two transfer functions to implement [2]. Following [65], $F(z)$ is normalised. Then the minimum phase zeros are selected [64] from $1 - F(z)F*(z)$ and used to construct the $H(z)$, the polynomial for the complementary output. In this way, power conservation and the complementarity of the outputs is guaranteed. The next step is the Jinguji recursive algorithm to find the coupling coefficients. As the expansion coefficients obtained from the Parks-McClellan method are symmetrical, the coupling coefficients are symmetrical with respect to the symmetry axis of the component. The delay length is calculated from the required FSR, and is equal for all stages.

As the filter spectrum is symmetrical, all expansion coefficients of $F(z)$ are real. Therefore, no phase shifts are needed in the filter stages. The filters area all linear phase.

The filter order N needed is plotted in figure 6.5 as a function of the number of channels and the required crosstalk level. Two different design points are shown, one where the half bandwidth is defined as $0.15\Delta f_{ch}$ at a level of -0.03 dB, and one where the 3 dB bandwidth is designed to be $2 \times 0.45\Delta f_{ch}$. The former was the design approach for the practical implementation, while the latter is more conventional. However, for a designed crosstalk of -30 dB, the results are similar. This result holds irrespective of channel spacing, as the filter order is defined by the requirement that N_{ch} channels fit in a *normalised* frequency range of 1 (the FSR). To reach a crosstalk level of -30 dB with a 3 dB bandwidth of $0.9\Delta f_{ch}$, the number of stages is approximately $N \approx 1.5 \cdot N_{ch}$. For a -50 dB crosstalk level, $N \approx 2.5 \cdot N_{ch}$. It is clear that only a limited number of stages will be possible for reasons of tolerance to fabrication variations[3], which limits the possible configurations. Limiting the 3 dB bandwidth to $0.6\Delta f_{ch}$ reduces the number of stages for a -30 dB crosstalk level to $N \approx 1.2N_{ch}$, so that 16 channel devices should be realizable. The tolerance to fabrication variations also depends on the delay length and thus the actual free spectral range, which will put another limit on the realisable configurations.

---

[2]In Matlab using the firpmord and firpm functions, in Python using the remez function in scipy and own library functions

[3]This is in contrast to digital filters, where a high order is easily achievable

Figure 6.5: Calculated filter order as a function of desired crosstalk, for different numbers of channels. Filters designed with the Parks-McClellan minimax approach. Solid line: 3 dB bandwidth = $0.9 \times \Delta f_{ch}$. Dotted line: 0.3 dB bandwidth = $0.3 \times \Delta f_{ch}$

6.2.4 Reducing phase error by using broader waveguides

Nanometre-scale width variations in the waveguides can cause phase errors. In lattice filters, these phase errors cause differences between the delay lengths, and so the total delay through a path becomes a different, non-integer multiple of the desired unit delay length for each path. This then generates higher sidelobes in the transmission spectrum and thus crosstalk.

Figures 2.10(b) and 2.13 in chapter 2 show the phase errors are considerably reduced for wider waveguides. The waveguides of delay lines in filters can be made broader, but bends should stay single mode. Therefore, the approach taken is to expand each narrow straight wire section by two tapers. This is illustrated by figure 6.6. The taper length for converting between 450 nm and 800 nm wide waveguides is chosen at a conservative 3 μm, which is short due to the high index contrast.

As broader waveguides have a different effective and group index, this needs to be taken into account when designing the filter for a given FSR or drop channel wavelength. As the group index becomes smaller with waveguide width, the delay length increases to keep the same free spectral range. This has an adverse effect on the phase error: going from 450 nm to

Figure 6.6: Broadening of delay lines for crosstalk reduction

800 nm wide wires makes the delay lines a factor 1.16 longer, and thus the phase error a factor 1.08 larger. However, the net effect is still a decrease of the phase error, as $\frac{d\beta}{dw}$ decreases with two thirds between those two widths. For high initial crosstalk, the reduced phase error can lead to a significant decrease of crosstalk with waveguide width. For low initial crosstalk, below $\sigma_{\delta\phi} = \frac{\pi}{500}\mu m^{-1/2}$, the gain is small, as will be shown later on. As the FSR is kept constant, the filter shifts over several filter orders when broader waveguides are used.

With CMOS fabrication methods, it is difficult to get both the narrow and wide waveguides on target at the same time, as discussed in chapter 3. Therefore, control over the exact wavelengths, which is already troublesome (see section 4.3.1), is traded for a reduction of the phase errors. By putting all the delay length in broader waveguides, this problem is circumvented, but then the couplers cannot be controlled at will. It is possible to compensate for the width differences by applying corrections to either waveguide, but extensive process characterisation is required for complex structures.

	N_{ch}	Δf_{ch} [GHz]	ripple [dB]	BW [GHz]	couplers	coupling coefficients (half)
1	4	400	0.2	30	6	0.09, 0.261, 0.418
2	8	400	0.2	30	14	0.0053, 0.0354, 0.0623, 0.1060, 0.1535, 0.1965, 0.2224
3	8	200	0.5	20	12	0.0323, 0.0606, 0.1047, 0.1549, 0.2005, 0.2282
4	4	200	0.3	25	6	0.0837, 0.2629, 0.4219
5	4	100	0.3	15	6	0.0731,0.2661,0.4285

Table 6.1: Specifications of designed filters with the Parks-McClellan/Jinguji method

6.2.5 Experimental results

In the designed filters, the number of stages (couplers) was limited in order to limit size and influence of fabrication variations. For example, the 11 stage filter from figure 6.4 has a footprint of $180 \times 40 \, \mu m^2$. For a maximum crosstalk of -30 dB designed, this limited the achievable bandwidth at ripple level to 15% - 25% of the channel spacing. Table 6.1 shows the designed filters (put on the PICMOS02 mask). All filters were implemented both with and without expanded waveguides, except for the 8 channel, 400GHz one.

Figures 6.7 and 6.8 show transmission measurements and simulated transmission spectra of the five lattice filters designed. Only the lattice filter in figures 6.7(c),6.7(d) was not implemented with expanded waveguides. The other measurements are for components with expanded waveguides. The FSR is reproduced exactly, and the 3 dB bandwidth deviates 10% maximally from the designed value, shown in table 6.2. Crosstalk is a lot higher than designed, however. With the expanded waveguides, sidelobes typically rise to -12 dB levels and down to less than -15 dB for some components. Table 6.3 lists the measured sidelobe level for each filter, both with and without expanded waveguides. For the longer total lengths (mainly for smaller FSR), the difference between expanded and normal width waveguides is significant and the broader waveguides deliver a gain of 3 to 5 dB, even though the delay lengths are longer to keep the same free spectral range. These crosstalk levels are too high for telecom applications. However, the broader waveguides deliver a low enough crosstalk for interconnect applications with port counts of those of the lattice filters studied here, as studied in chapter 9.

(a) 4×400 GHz,simulated

(b) 4×400 GHz,measured

(c) 8×400 GHz,simulated

(d) 8×400 GHz,measured

(e) 8×200 GHz,simulated

(f) 8×200 GHz,measured
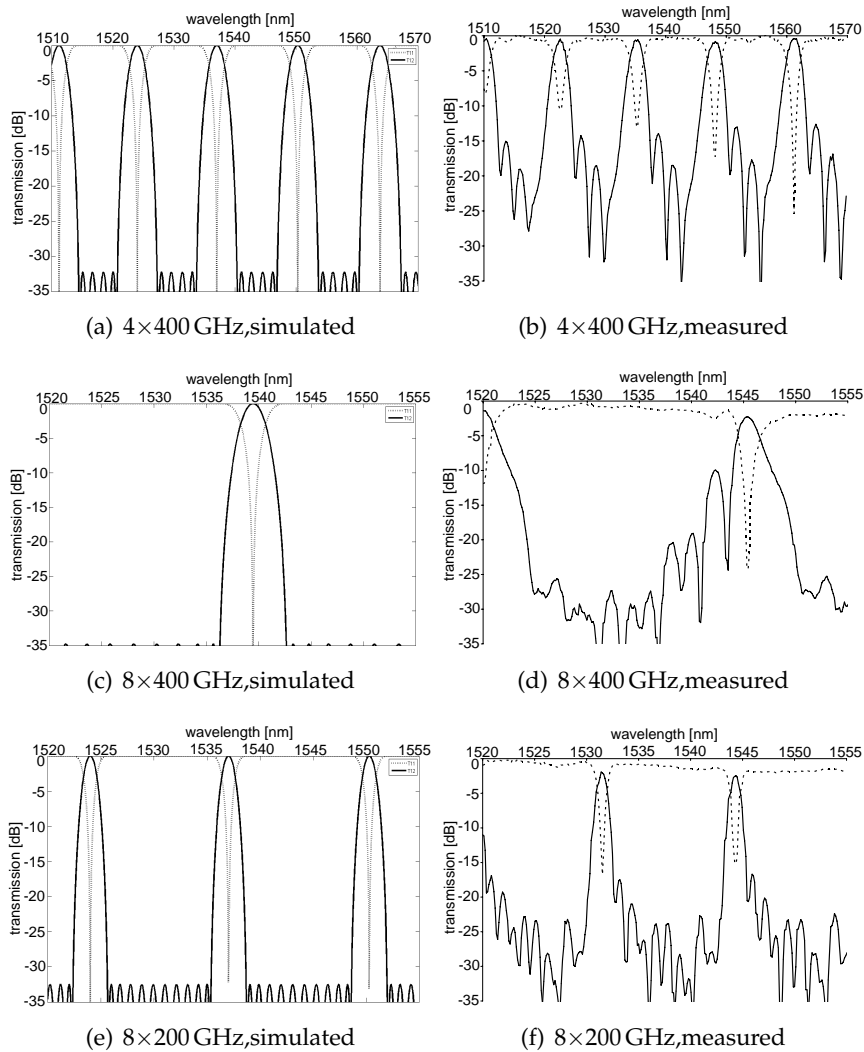
Figure 6.7: Simulated and measured transmission spectra of lattice filters (part 1). The device in 6.7(c), 6.7(d) did not have expanded waveguides, while the others did have.

(a) 4×200 GHz, simulated

(b) 4×200 GHz, measured

(c) 4×100 GHz, simulated

(d) 4×100 GHz, measured
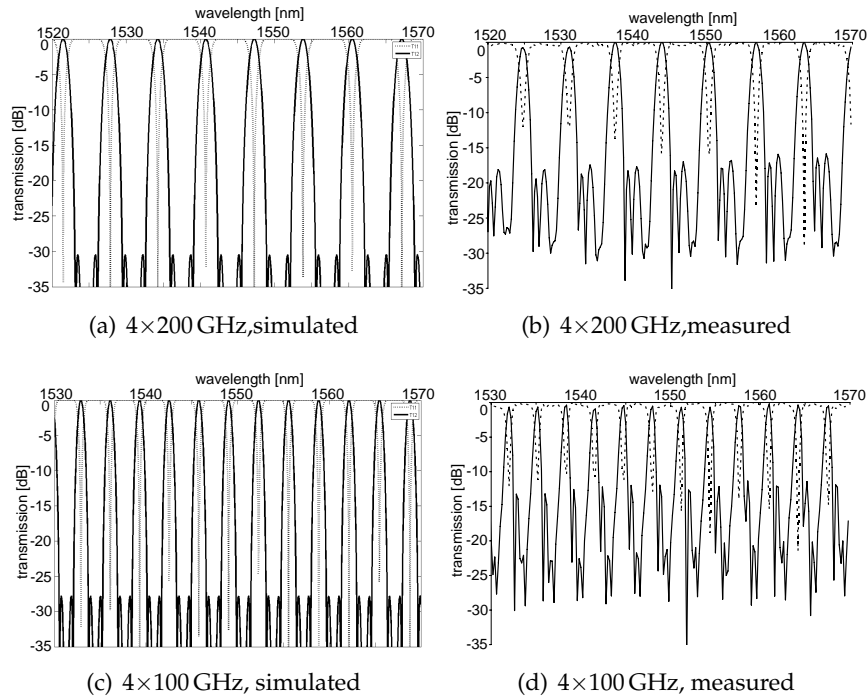
Figure 6.8: Simulated and measured transmission spectra of lattice filters. The device in 6.7(c), 6.7(d) did not have expanded waveguides, while the others did have.

	designed BW	measured BW
1	2.55 nm	2.28 nm
2	2.36 nm	2.18 nm
3	1.25 nm	1.15 nm
4	1.3 nm	1.2 nm
5	0.66 nm	0.62 nm

Table 6.2: 3 dB bandwidth of the designed and measured lattice filters

	FSR [GHz]	N	exp. waveg.	ΔL [µm]	total length [µm]	mean CT_{max} [dB]	$\sigma_{CT_{max}}$ [dB]
1	1600	5	no	41.52	207.6	-12.8	4.0
		5	yes	46.125	230.625	-11.7	2.4
2	3200	13	no	20.76	269.88	-7.9	2.0
3	1600	11	no	41.52	456.72	-5.0	2.9
		11	yes	46.125	507.375	-10.7	2.4
4	800	5	no	83.04	415.2	-11.4	2.0
		5	yes	93.715	468.575	-10.7	2.3
5	400	5	no	166.08	830.4	-7.90	1.8
		5	yes	188.89	944.45	-11.8	2.9

Table 6.3: Measured FSR, delay length and sidelobe (CT) level of the lattice filters. Six nominally identical samples were measured for each filter. The isolated narrow wire width was 480 nm, the expanded wire width 840 nm.

6.2.6 Sources of crosstalk

To extract data on the fabrication variations causing the crosstalk, the influence of the different contributions is simulated in the next sections. The possible contributions are

- **Deterministic phase noise in the waveguides** due to mask digitisation.

- **Stochastic phase noise in the waveguides** due to waveguide width variations caused by the mask, lithography and etching, and due to silicon top layer thickness variations.

- **Deterministic power coupling variations** in the directional couplers due to a deterministic difference of the designed waveguide widths and heights, the gap and a non-rectangular form factor of the waveguide core.

- **Stochastic power coupling variations** in the directional couplers due to similar but stochastic effects caused by the mask, lithographic irregularities, etching and silicon top layer variations.

- **Deterministic and stochastic phase errors in the directional couplers.** Due to non-ideal fabrication and the intrinsic problem of exciting a high index contrast directional coupler, small stochastic geometrical variations not only influence the power coupling of the

directional coupler, but also caused the outputs of the couplers not to be in quadrature.

It is difficult to separate the deterministic and stochastic phase noise contributions in the waveguides, but all filters have been designed with their straight waveguides aligned with the mask grid. However, the small bends may suffer mask digitisation effects.

The influence of deterministic and stochastic coupler variations is somewhat different from that of phase errors. This can be understood from the Z transform description. A phase error changes both the angle and magnitude of several zeros. Evaluated on the unit circle, this translates in an alteration of the optical frequency corresponding to the affected zeros and the depth of the transmission dips. This also changes the position and height of the sidelobes, and breaks the symmetry of the transmission spectrum. The effect of phase errors is most clear on the ripple in the stopband of the drop port. The zeros are optimised for the lowest possible equiripple, phase errors will make the sidelobes in the stopband of the drop port deviate from the equiripple behaviour and will increase the height of some of the sidelobes. The higher sidelobes lead to insertion loss in the pass port. The influence on the extinction in the pass port on the drop channel is generally more limited, and only becomes appreciable for high phase errors.

On the other hand, a coupling error alters the magnitude of several zeros only, moving them closer to or further away from the unit circle. As optical transmission is evaluated on the unit circle, this has an influence on the depth of the transmission dips defined by the zeros. The effect is a reduction of extinction in the stopband of the pass port, and the increase of the ensemble mean of the sidelobe level in the stopband of the drop port. Of course, the effect of coupler variation is only depending on the filter order and the position of the actual zeros, and not on the free spectral range or channel spacing as with phase errors.

Normally, a filter is designed so that the zero frequencies coincide with the frequency channels neighbouring the central channel, and thus it is the change of the extinction in the transmission dips which has the largest influence.

For random errors, the variance of the sidelobe level and the extinction will of course increase with the variance of the phase or coupler errors.

Note that a single phase or coupler change will influence the location of several zeros at once, which is exactly why lattice filters are difficult to be designed by hand, but also why they are so flexible. Both influences are illustrated using a simple second order filter, plotted in figure 6.9.

(a) $\frac{1}{3} + \frac{1}{3}z^{-1} + \frac{1}{3}z^{-2}$

(b) $0.3333 + 0.3333z^{-1} + (0.3170 - 0.1030i)z^{-2}$ ($\frac{\pi}{10}$ phase change of z^{-2})

(c) $0.3676 + 0.3557z^{-1} + 0.2767z^{-2}$
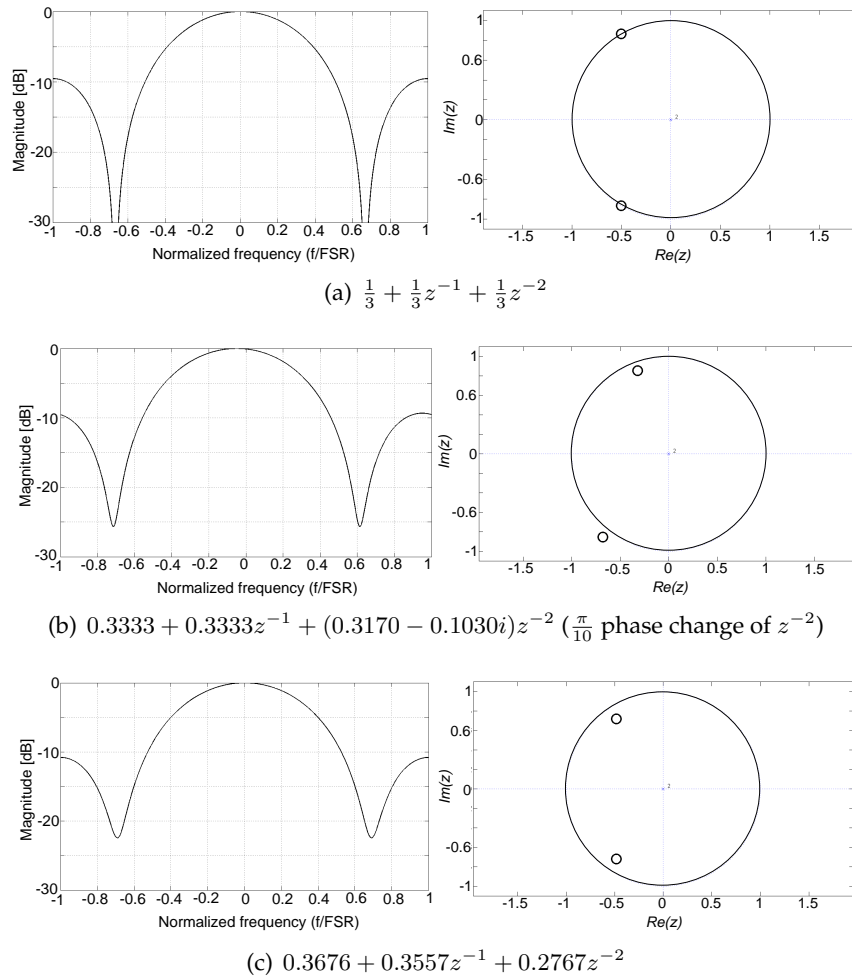
Figure 6.9: Simulated influence of phase and coupler deviations on the theoretical transmission and zeros of a second order lattice. (a) original filter (b) with phase error (c) with coupler error

All these effects can be simulated, either by taking random coupler variations into account or by varying the coupling κ between the waveguides. The two effects can be combined in order to see the aggregate effect.

6.2.7   Phase noise in the waveguides

To some extent, the higher than designed sidelobes are due to phase deviations in the waveguides. Based on the simple model laid out in chapter 2, we can get an order of magnitude feeling of the influence of phase noise on lattice filters.

Most channel drop filters will have a delay length of at most 200 µm, as systems either use a limited number of channels with 100 or 200 GHz channel spacings, or a larger number of channels with 25 or 50 GHz spacing. Therefore, the phase noise in each stage is limited. However, the zero locations are defined by all stages together, and thus the maximum path length becomes important. Thus, the higher the filter order, the more important phase errors become. In contrast, with cascades of Mach-Zehnders as in section 6.3, the zero locations are controlled by each stage independently and phase errors would only have an influence per stage[4].

For interleavers, phase errors are more important, as a 100 GHz channel spacing or smaller requires a delay length of more than 650 micrometre.

With a phase error that can easily reach $\sigma_{\delta\phi}/\sqrt{L} = \frac{\pi}{200}\mu m^{-1/2}$ to $\frac{\pi}{100}\mu m^{-1/2}$ for $\sigma_d = 5\,\text{nm}$ and $L_c = 500\,\text{nm}$, a unit delay line can experience a significant $\sigma_{\delta\phi}$ of $\frac{\pi}{14}$ to $\frac{\pi}{7}$.

Figure 6.10 shows a simulation of the transmission spectra of the 4x100 GHz lattice filters, with random phase errors in each delay length. It is clear the performance can be seriously deteriorated for sidewall variation with a standard deviation of 5 nm and a submicrometre correlation length.

The performance can be analysed in more detail by looking at the simulated actual crosstalk as a function of designed crosstalk (filter order), for a given standard deviation of phase error. Enough ensemble members need to be calculated to get a representative result. In the following, 50 samples were taken for each parameter and we look at the mean and peak values of the crosstalk, where of course one should remember that the distribution tail extends beyond the calculated extremal samples. The phase error is considered to be wavelength independent, so $\frac{d\beta}{dw}(\lambda)$ is constant.

---

[4]However, the sidelobes or absence of extinction of each stage will add up

(a) Original design

(b) With phase errors, $\sigma_{\delta\phi}/\sqrt{L} = \frac{\pi}{200}/\sqrt{\mu m}$, 10 samples
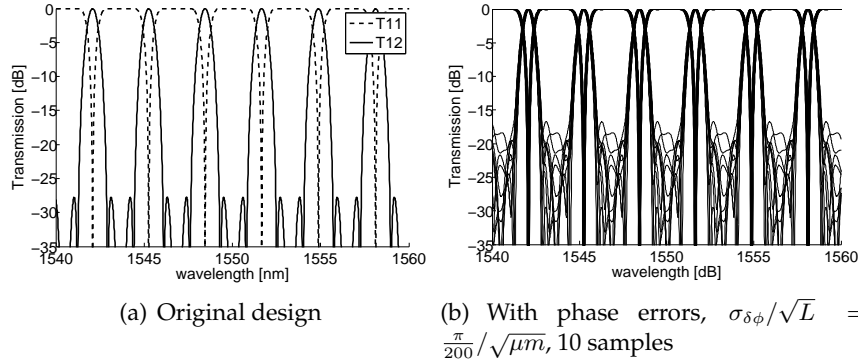
Figure 6.10: Simulated transmission spectra of lattice with 4x100GHz channels with phase errors

First a four-channel, 100 GHz system is considered, as this is a nearly worst case scenario. Simulation results are given in figure 6.11. For a given standard deviation of the phase error per square root micrometre, the standard deviation of the generated crosstalk clearly grows with the filter order (lower designed crosstalk). With a phase error of $\pi/400\,\mu m^{-1/2}$, the mean crosstalk to expect for -30 dB designed is already -25 dB, peaking at around -19 dB. The mean insertion loss is below 1 dB. Adding stages reduces the mean crosstalk with less than 5 dB. With a phase error of $\pi/200\,\mu m^{-1/2}$, mean and peak crosstalk are 5 dB higher, at -20 dB and -15 dB, respectively. Increasing the phase error to $\pi/50\,\mu m^{-1/2}$ increases the crosstalk with another couple of dB, and in extreme cases the filter characteristic is completely destroyed.

To have a better look at the influence of the number of stages, the channel spacing and the designed crosstalk, the average expected crosstalk and the standard deviation were calculated from 50 samples with stochastic phase errors for each filter. The result for a low and a high level of phase errors and for both -50 and -30 dB designed crosstalk is given in figure 6.12. Increasing the number of channels while keeping the spacing constant increases the FSR and thus shrinks the delay length, but requires more stages, so the influence of the number of channels on the expected crosstalk is small. This is an interesting point with respect to scalability, as increasing the number of channels does not put a tighter requirement on the phase error. However, with a higher channel count, a lower crosstalk level is needed, so that the achievable number of channels with current technology is still limited. On the other hand, the channel spacing has

(a) $\sigma_{\delta\phi} = 0.0$

(b) $\sigma_{\delta\phi} = \frac{\pi}{5000}\mu m^{-1/2}$

(c) $\sigma_{\delta\phi} = \frac{\pi}{400}\mu m^{-1/2}$

(d) $\sigma_{\delta\phi} = \frac{\pi}{200}\mu m^{-1/2}$

(e) $\sigma_{\delta\phi} = \frac{\pi}{100}\mu m^{-1/2}$

(f) $\sigma_{\delta\phi} = \frac{\pi}{50}\mu m^{-1/2}$
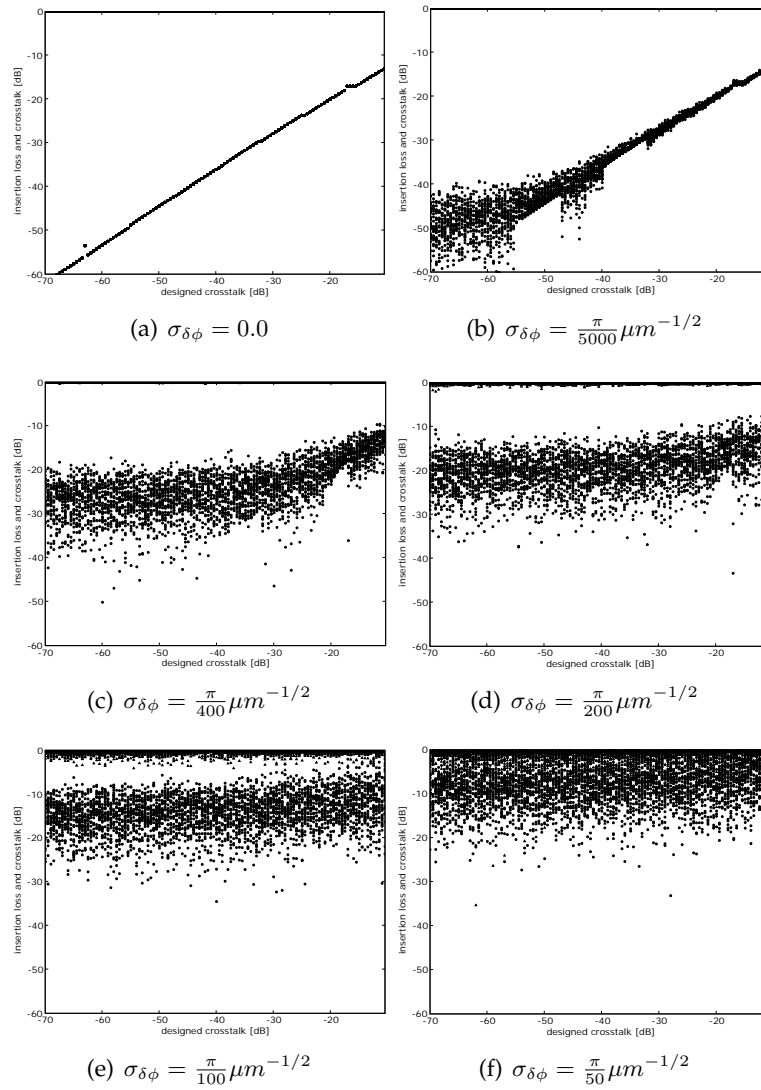
Figure 6.11: Simulated insertion loss and crosstalk of a lattice filter with 4 channels and $\Delta f_{ch} =$ 100 GHz, as a function of designed crosstalk, for different phase error standard deviation

(a) $\sigma_{\delta\phi} = \frac{\pi}{1000}\mu m^{-1/2}$, CT=-30 dB designed (b) $\sigma_{\delta\phi} = \frac{\pi}{1000}\mu m^{-1/2}$, CT=-50 dB designed

(c) $\sigma_{\delta\phi} = \frac{\pi}{100}\mu m^{-1/2}$, CT=-30 dB designed (d) $\sigma_{\delta\phi} = \frac{\pi}{100}\mu m^{-1/2}$, CT=-50 dB designed
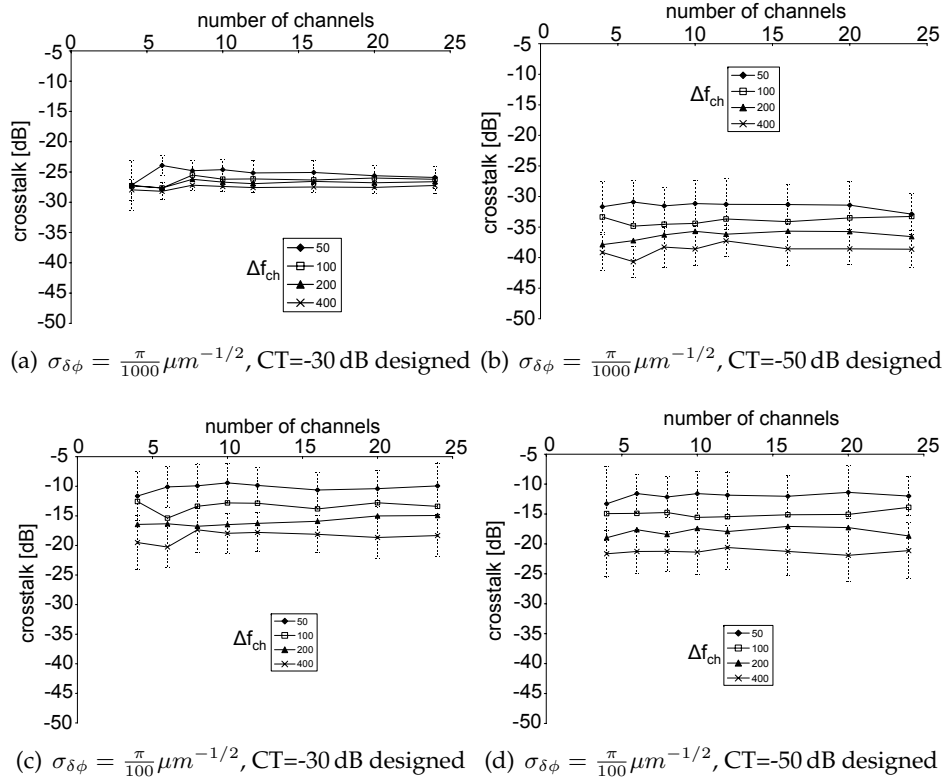
Figure 6.12: Calculated mean crosstalk of lattice filters as a function of channel count and channel spacing, for two phase error levels and -30 dB and -50 dB designed crosstalk. For 50 GHz and 400 GHz channel spacing, the standard deviation is plotted as error bars

a large influence for high phase errors or a low designed crosstalk (high filter order). For a phase error of $\pi/100\,\mu m^{-1/2}$, the crosstalk resulting from phase errors reduced by almost 10 dB when going from 50 GHz to 400 GHz channel spacings. Designing the filters for lower crosstalk of -50 dB does not help much for relatively high phase error levels. And even at a low phase error level, the expected crosstalk stays at least 10 dB above this level.

From these figures, the measured transmission spectra of the lattice filters could be explained by a phase error of the order of $\frac{\pi}{100}$ to $\frac{\pi}{50}$. However, the reduced extinction in the stopband of the pass port, and the unclear dependency of measured crosstalk on the channel spacing indicate coupler variations have a large contribution to crosstalk in the lattice filters.

6.2.8   Coupler variations

In the couplers, both the power and the phase relation between the outputs can change. The latter effect adds up to the phase errors in the waveguides, while the former reduces the extinction of the transmission dips and alters the sidelobes. Normally, the phase between two outputs of a directional coupler is $\frac{\pi}{2}$, but this can change due to losses in the coupler. The coupler variation can have a deterministic and a stochastic component.

In the designs, the coupling κ was kept constant at $0.04/\mu m$ (constant gap and waveguide widths) and the power coupling was varied by changing the coupler length. The influence of coupler variations is illustrated by figure 6.13, which gives the influence of a deterministic variation and stochastic coupler variations on the lattice filter with four 100 GHz spaced channels.

If the gap or waveguides are narrower or wider than intended due to a different exposure dose during lithography, or the silicon thickness is different than taken into account during design, all couplers will couple less or more. The effect is not linear in function of designed coupling, as the dependency of coupled power on the widths and gaps is complex. However, the behaviour is deterministic. As a simplification, in figure 6.13 a deterministic coupling deviation that scales linearly with coupling coefficient is supposed. The 10% and 30% deviations reduce the extinction for the drop channel considerably, and for a 30% deviation, the passband loss of the drop channel is also considerably increased.

A stochastic coupler variation has its most visible influence on the stopbands of the drop port, where the sidelobes are increased with 10 dB to 20 dB in the example, and in the stopband of the pass port, where the extinction is also reduced.

6.2.9   Combining random phase errors and coupler variations

The influence of simultaneous random variations of the directional coupler coupling strengths and phase errors in the waveguides were simulated for each of the fabricated devices. Figures 6.14 and 6.15 plot the estimated mean crosstalk to the drop port, computed with 100 random samples for each point, with $\sigma_{\delta\phi}$ ranging between $\frac{\pi}{1000}\mu m^{-1/2}$ to $\frac{\pi}{20}\mu m^{-1/2}$ and $\sigma_{\delta\kappa}$ between 0.003 and 0.3, where the influence is highest. The extinction was calculated at the channel wavelengths, so that this is the actual crosstalk seen. Part of the irregularity is due to the limited accuracy of the estimation of the mean crosstalk. Due to the complex nature of the zero locations as a function of coupling and phase values, the resulting

(a) constant deviation 10%

(b) constant deviation 30%

(c) random deviations: Gaussian, $\sigma_{\delta\kappa}$=10%, 10 samples

(d) random deviations: Gaussian, $\sigma_{\delta\kappa}$=20%, 10 samples
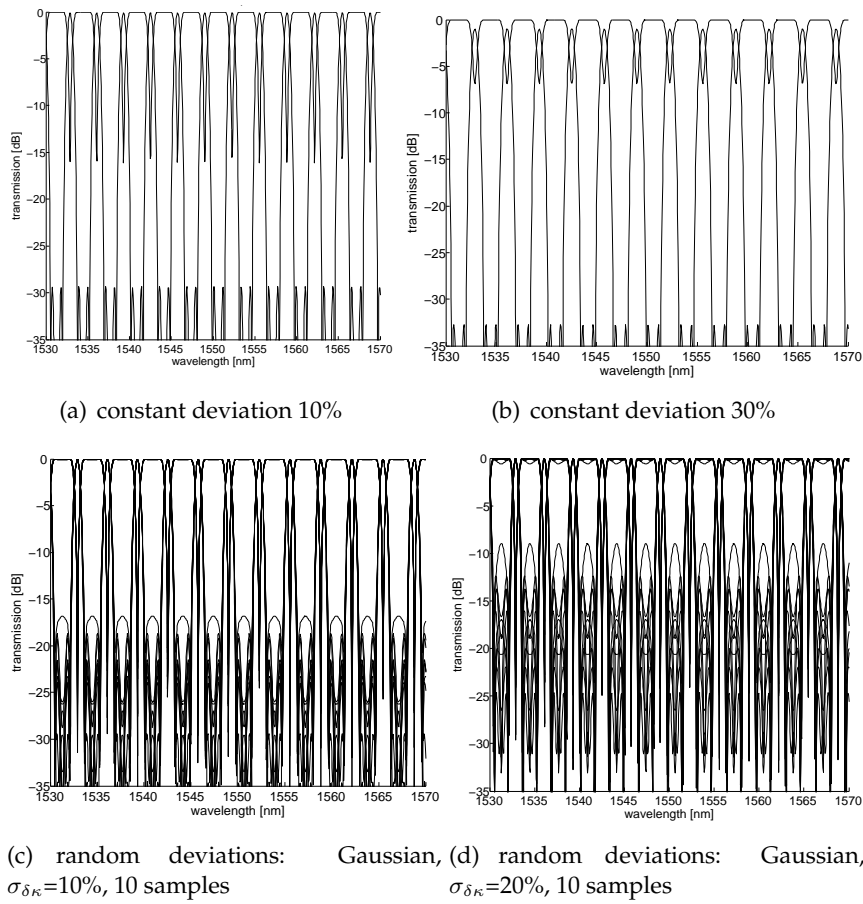
Figure 6.13: Calculation of influence of constant deterministic and random coupler variations on the transmission spectrum of a 4 channel lattice with 100 GHz spacing.

(a) $4 \times 100\,GHz$, no expansion

(b) $4 \times 100\,GHz$, with expansion

(c) $4 \times 200\,GHz$, no expansion

(d) $4 \times 400\,GHz$, no expansion

Figure 6.14: Influence of random coupler and phase variations on the PICMOS02 lattice filters: calculated mean extinction in the drop port at the channel wavelengths. (part 1)

crosstalk from both effects is not additive. The error of simply adding the generated crosstalk was calculated to be up to a couple of dB, which is important for high crosstalk. The larger crosstalk for smaller FSR due to phase errors is clearly visible, as well as the larger crosstalk for a larger number of channels due to phase and coupler errors.

With expanded waveguides, the crosstalk is slightly larger for the same level of error due to the longer delay length. However, the net effect is a reduced crosstalk due to a reduced phase error. Moving from $\frac{\pi}{30}$ to $\frac{\pi}{120}$ would reduce the crosstalk from -5 dB to -15 dB. For smaller initial phase error, the gain is smaller.

As important is the extinction of the dropped channel in the pass port, which decreases greatly if phase and coupler errors are combined. This

(a) $8 \times 200\,GHz$, no expansion

(b) $8 \times 200\,GHz$, with expansion

(c) $8 \times 400\,GHz$, no expansion
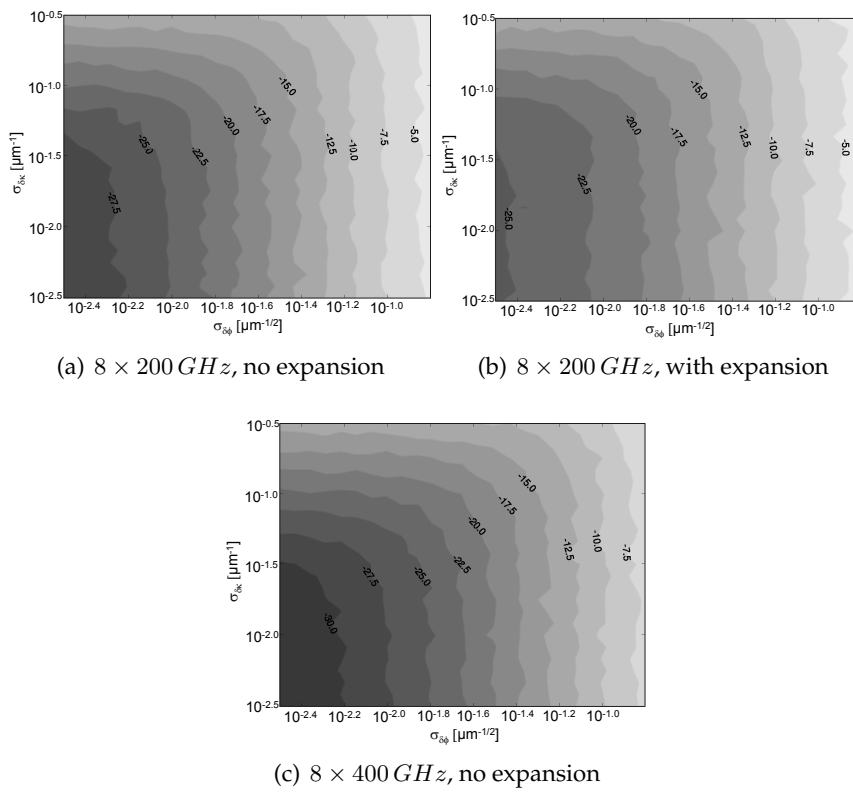
Figure 6.15: Influence of random coupler and phase variations on the PICMOS02 lattice filters: calculated mean extinction in the drop port at the channel wavelengths. (part 2)

(a) $4 \times 100\,GHz$, with expansion

(b) $8 \times 400\,GHz$, with expansion

Figure 6.16: Influence of random coupler and phase variations on the PICMOS02 lattice filters: calculated mean extinction in the pass port at the drop channel wavelength

can be seen from figure 6.16, which shows the extinction for the 4x100 GHz and 8x400 GHz filters. For serious phase errors, only a very high coupling error of 15% to 30% decreases the extinction even more. The extinction can reach values as low as 5 to 10 dB.

6.2.10 Interpretation of the measurements

Coupler errors or phase errors alone cannot explain the measured crosstalk level. Only a large phase error pushes the crosstalk to the measured level of -5 dB to -10 dB for the lattice filters. From the figures, the crosstalk level of -5 dB measured for lattice filters without expanded waveguides can be explained by a large phase error of the order of $\frac{\pi}{30}\mu m^{-1/2}$, almost irrespective of the coupler error. Note that part of the phase error may be related to the couplers and not to the waveguides. The two contributions cannot be separated here. Still, in order to explain the level of phase errors through waveguide phase error only, the sidewall variations should have a σ_d of around 10 nm for a correlation length of a few 100 nm to a micrometre, corresponding with width variations with a σ of 14 nm. Given the SEM pictures in chapter 3 and the uniformity measurements in chapter 8 later on, this seems rather large. Therefore, the crosstalk levels here indicate that in addition to short- and midrange phase errors, other effects are present: stochastic coupler phase variations, stochastic mean linewidth variations on the mask, or, most likely, high coupled power variations.

From figure 6.16, we can see in the pass port a drop channel extinction of 15 to 20 dB should be feasible, which is confirmed by the measurements (figures 6.7,6.8). From the measurements, we can also see smaller extinction, specifically for shorter wavelength. This again indicates strong coupler variations from the designed values, superimposed on the phase errors. Given the wavelength dependency of the directional couplers, this is plausible.

The lattice filters with four channels and expanded waveguides show a crosstalk signature that is very similar for 6 nominally identical samples. Figure 6.17 plots their cross port transmission spectra. The passband locations varied over a few nm, but were aligned so that the figure is more clear. The zero frequencies have a large offset from their ideal positions, or there are even less zeros visually clear than expected. The offset reaches $\frac{\pi}{5}$ (with 2π the FSR), but the real zero positions are remarkably similar for most samples of each filter, with standard deviations of the order of $\frac{\pi}{10}$ to $\frac{\pi}{60}$. This indicates an important deterministic contribution to phase noise, either caused by waveguide phase error or by non-ideal phase transfer of the directional couplers. One possibility is differences in mask digitisation between the waveguide bends, as all straight lines are oriented along the grid. However, more plausible is a deterministic term in the phase transfer of the couplers, deviating from $\frac{\pi}{2}$ but depending on the length of the coupler.

The variation of the zeros frequencies around their mean is then due to the stochastic component in the waveguide phase error and coupler phase deviation. The differences in zero depth can additionally be caused by power coupling deviations, and the differences in sidelobe height are caused by the deviations in zero locations, and by the phase errors and power coupling errors as such. For all three filters with expanded waveguides, the mean maximum sidelobe height is between -10.7 dB and -11.8 dB, and the standard deviation of maximum sidelobe height is 2.3 dB to 2.9 dB. From the indicated channel positions, we can see the channels do not coincide approximately with the zeros as designed, but with the sidelobes, giving rise to significant crosstalk which is thus around -11 dB maximum. For the lattice filters with 8 channels and expanded waveguides, the same conclusion can be made, but the spectrum is more complicated as the order is higher. An example with 6 nominally identical samples is given in figure 6.17 too. Its highest sidelobe has a mean level of -10.7 dB, with a standard deviation of 2.4 dB. The nearest-neighbour crosstalk is as high as -5 dB, while the other channels see a crosstalk lower than -15 dB.

(a) $4 \times 100\,GHz$

(b) $4 \times 200\,GHz$

(c) $4 \times 400\,GHz$

(d) $8 \times 200\,GHz$
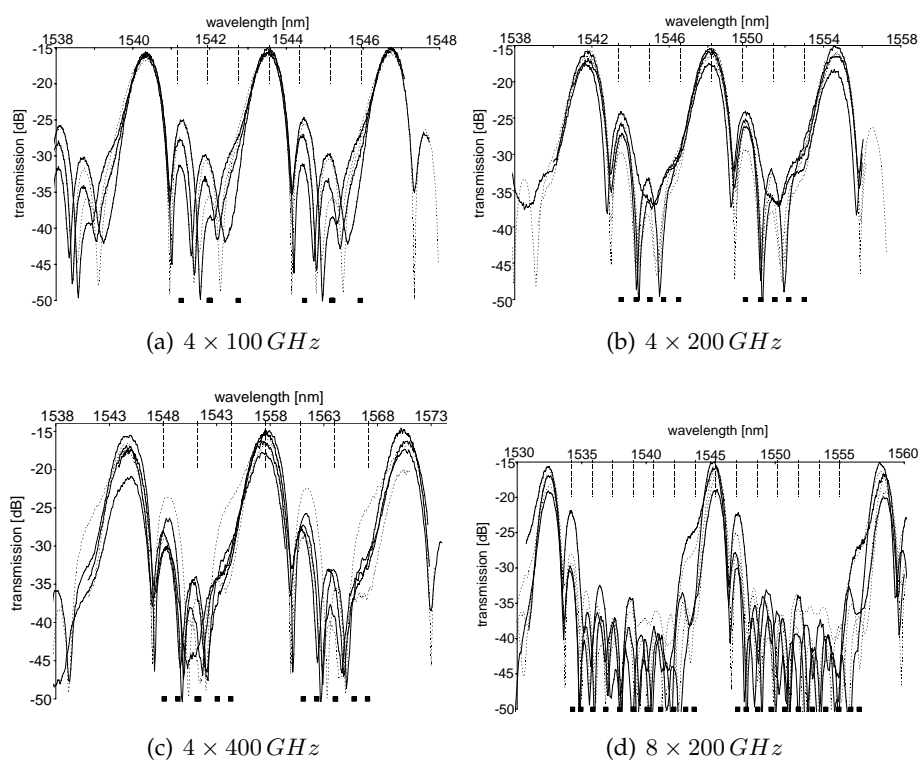
Figure 6.17: Measured transmission spectra of lattice filters with expanded waveguides on 6 nominally identical samples, with the passbands aligned. Black squares indicate the designed zero locations, broken lines at the top indicate channel locations.

When the delay lines are not broadened, analysis becomes more difficult as the distortion is generally higher. A thorough analysis and comparison of filters with and without waveguide broadening is difficult, as the low number of samples limits the accuracy of the mean and standard deviation estimates. For instance, the samples of the $4 \times 400\,GHz$ filter without expanded waveguides have a mean highest sidelobe level of $-12.8 \pm 3.97\,$dB, which is lower but with a higher σ than those of the equivalent filter with expanded waveguides. The smaller gain apparent from the characterised filters with expanded waveguides also points either to a large coupling error, of the order of $\sigma_{\delta\kappa} = 0.2$ to 0.3, or to a coupler-induced phase error, or, most presumably, both.

6.2.11 Predicting performance

Requirements for a better performance can be deduced from figures 6.14, 6.15. To get the crosstalk consistently under the -20 dB mark, also for 100 GHz channel spacing, the phase errors should be reduced from the order of $\frac{\pi}{50}\mu m^{-1/2}$ to $\frac{\pi}{300}\,\mu m^{-1/2}$, and random coupler variations should have a $\sigma_{\delta\kappa}$ below $0.1/\mu m$. For smaller channel spacings of 50 GHz and below, the requirements are more stringent. The phase errors in the delay lines can be reduced by reducing their amplitude and their correlation length. Deviations along the wire on a length scale of 500 nm to a few micrometre should be removed, or their amplitude should be reduced to 2 to 3 nm, but then broadened waveguides are still needed. Additionally, the deterministic and stochastic coupler errors should be reduced. In case only shorter length scale variations remain, waveguide broadening is probably not really needed for -20 dB crosstalk. To obtain a crosstalk better than -30 dB, the phase errors need to be lower than $\frac{\pi}{1000}\,\mu m^{-1/2}$ and the filters need to be designed for -50 dB crosstalk at the expense of more stages and thus larger footprint. However, these should scale to 16 and 24 channels easily. For longer correlation lengths, this means the roughness amplitude needs to be far below 1 nm, even with broader waveguides. But even for shorter correlation length, the standard deviation of roughness should not exceed 1 nm.

What performance can we expect in the near future ? The main improvement would be to remove the deviations on a length scale of 500 nm to a few micrometre. However, if their origin is already on the reticle, as we will see in chapter 8 multiple sources and measurement results in this work suggest, then we can only hope to reduce their amplitude through processing. With σ_d reduced to 2 to 3 nm, phase errors should be reduced to $\frac{\pi}{600}$ to $\frac{\pi}{500}\,\mu m^{-1/2}$. Additionally, the directional couplers

should be better controlled, with respect to both power and phase, which needs larger process windows and extensive characterisation. Designing for lower crosstalk can win a few dB in crosstalk, but will not make a huge difference at these phase error levels.

Therefore, in the relatively near future, through better processing it should be possible to push the crosstalk consistently below -20 dB, with a small tail of the component distribution above -20 dB. This will work out most easily for channel spacings of 200 and 400 GHz, but should be possible for 100 GHz spacing too, irrespective of the number of channels. With a -20 dB crosstalk, 16 channel operation should be feasible.

A lower crosstalk of the order of -30 dB will only be gained through drastically better fabrication technology or through tuning of the individual waveguides, for instance using heaters.

6.3   Mach-Zehnder cascade

Cascading MZIs is another way to build higher-order channel selectors or (de)multiplexers. In this configuration, the zero of each stage can be set independently, and thus easier design is possible. A method to implement channel drop filters or even demultiplexers is to cascade MZIs with each MZI having double the FSR of its predecessor. In this way, the component scales only logarithmically with the number of channels, and still has relatively low passband loss [64]. However, there is less room for tailoring the frequency response than is possible with lattice filters. They are not ideal as channel selectors as the spare outputs need to be combined back, resulting in excess loss. For low channel counts however, a MZI cascade demultiplexer can be more space efficient than other filters such as arrayed waveguide gratings.

Phase errors in the waveguides of a stage displace the zero location of that stage, and therefore have an influence on the transmission of the stage the next stages, but does not influence the zero locations of the next stages as with a lattice filter. Therefore, the influence of phase errors may be smaller than with lattice filters. As all directional couplers are the same, ideally 3 dB splitters, there is no deterministic variation of the coupler phase transfer and corresponding zero location offsets. However, stochastic variations have the same influence as waveguide phase errors. In contrast to lattice filters, extinction and a lack thereof are directly transferred to the next stage, and the extinction values add up.

(a) Schematic

(b) SEM picture

Figure 6.18: Test structure for a Mach-Zehnder cascade

6.3.1   Measurements

Mach-Zehnder cascade channel selectors were designed with 100, 200 and 400 GHz channel spacings. Three MZIs were cascaded in one structure, monitoring both outputs of the last MZI and the spare outputs of the first two MZIs. This is illustrated in figure 6.18. In this way, a 1-out-of-2 channel selection (interleaver), 1 out of 4 and two 1 out of 8 channel selecting outputs could be tested in one structure.

Figure 6.19(a) plots the transmission to the four outputs for a 400 GHz channel spacing. The filter is functionally doing what it should, but the lack of extinction is apparent. The 1 out of 2 transmission is the *bar* transmission of the first MZI, with an extinction of 7 to 15 dB. In the cross port, to which the second and third stage are connected, the extinction is much higher, as seen in section 6.1.2, which is clear from the sharp zeros in the spectrum of the further output ports. The sidelobes of the 1 out of 4 port (bar transmission of the second MZI) are only at -5 dB of the main lobe, as the zero at that frequency is to weak to be even visible. The sidelobes in the outputs of the third stage are at about -5 dB from their main lobes too, mainly caused by the lack of extinction in the first stage, while the too small extinction in the second stage lets the more minor sidelobes rise. The simulated transmission with a power coupling value of 30% and 4% loss in the couplers is given in figure 6.19(b). Comparing the measured spec-

trum, the zero in transmission in the 1/8 outputs under the 1/4 main lobe is displaced from its designed position (indicated by the circle), originating from a phase error of about $\frac{\pi}{6}$ in the second stage, and the zero under one of the 1/8 outputs is missing completely.

6.4 Conclusion

In this chapter, we saw that Mach-Zehnder lattice channel drop filters are already feasible with CMOS fabrication technology for 100 GHz and higher WDM channel spacings. However, crosstalk is still around -10 dB and the yield is limited. This currently limits the scalability of the filters to 4 to 8 ports. Nevertheless, these filters fit in an area as small as 0.01 mm².

In chapter 9, we will see even at such a performance real application is already possible, albeit with a limited channel count. Mach-Zehnder cascades where fabricated but still showed high crosstalk, reflecting the relatively low extinction of the individual Mach-Zehnder structures.

Further enhancement of these structures mainly needs better directional couplers, and lower phase deviations in delay lengths due to roughness. With some process development, it should be possible to gain enough in crosstalk to enhance the scalability with a factor of two. To obtain lattice filters with drastically lower crosstalk, of the order of -30 dB, a drastically better fabrication process is needed, or tuning of each individual waveguide. Fabrication process improvement is necessary anyway in order to decrease the complexity of a tuning algorithm, by making components more reproducible and predictable.

(a) Measurement

(b) Simulation

Figure 6.19: Transmission spectra of logarithmic MZI cascade test structure with $\Delta f_{ch} = 400\,GHz$. (a) Measurement (b) Simulation, with 30%:70% coupling balance instead of 3dB couplers, and 4% (0.2 dB) loss in the couplers.

Chapter 7

Arrayed waveguide gratings

The arrayed waveguide grating (AWG) is about the only integrated passive wavelength filtering device actually commercially available on a relevant scale. Its elegant concept supports demultiplexing and routing of many parallel wavelength channels. It consists of an array of waveguides implementing the dispersive function, and two star couplers for dispersively focussing of input and output beams. Just as with ring resonators and Mach-Zehnder based filters, a high index contrast enables AWG devices of extremely compact size, with footprints of a fraction of a square mm.

In this work, AWG devices in silicon-on-insulator were simulated, fabricated and characterized. An analysis of the loss and crosstalk was performed, and the influence of the governing parameters for devices in SOI was studied experimentally.

First a brief overview of the general properties of the AWG is given and then the intrinsics of AWG in SOI are discussed. Further, simulation of AWG devices is discussed and the performance of fabricated devices is reported on. This gives a good idea of the current performance of components fabricated with CMOS-based technology. As in the previous chapter, we focus on the origins of the crosstalk and how to improve on it.

The AWG is also called phased array (PHASAR) or waveguide grating router (WGR), which mostly depends on the research group and how the AWG is seen and reasoned upon.

7.1 Basic properties

The AWG consists of an array of waveguides, typically with a constant length increment between them, acting similar to a phased array antenna. At the fan-in and fan-out, star couplers connect the input and output wave-

Figure 7.1: Schematic of an arrayed waveguide grating (AWG) with N_{in} input ports, N_{out} output ports, and N arrayed waveguides with a length increment of ΔL.

guides with the waveguide array. These star couplers are frequently called the free propagation regions. This is illustrated by figure 7.1.

The AWG acts similar to a concave grating, with the arrayed waveguides periodically mounted on the grating line and the input and output waveguides mounted on the focus lines. It implements both a dispersive function and a focussing function. The wavelength-dependent phase difference between the apertures on the grating line is controlled by a length difference between the arrayed waveguides, and light is diffracted open and focused by two star couplers. Those two components can be independently designed to some extent.

Given its similarity to concave gratings, general grating theory can be used to reason about AWG devices. From the optical path function[1], the grating equation is derived by setting first the derivative equal to zero (Fermat's principle), and the aberrations on the focus are given by the

---

[1]The optical path function can be written as [75]

$$F\left(w\right) = n_s r_i\left(w\right) + n_g L\left(w\right) + G\left(w\right) m'\lambda + n_s r_j\left(w\right) \tag{7.1}$$

with n_s and n_g effective indices in the slab and waveguides respectively, $L\left(w\right)$ the length of the waveguide between (u, w) and (u', w'), $G\left(w\right)$ the groove function being the number of waveguides counting from the center, and r_i and r_j are the distance from w to i and j. Both $L\left(w\right)$ and $G\left(w\right)$ are extended to a continuous function. The $r\left(w\right)$ and $L\left(w\right)$ terms mainly determine the focussing operation of the AWG, while the $G\left(w\right)m'\lambda$ term mainly determines the dispersive grating diffraction.

(a) Rowland mounting   (b) Confocal mounting
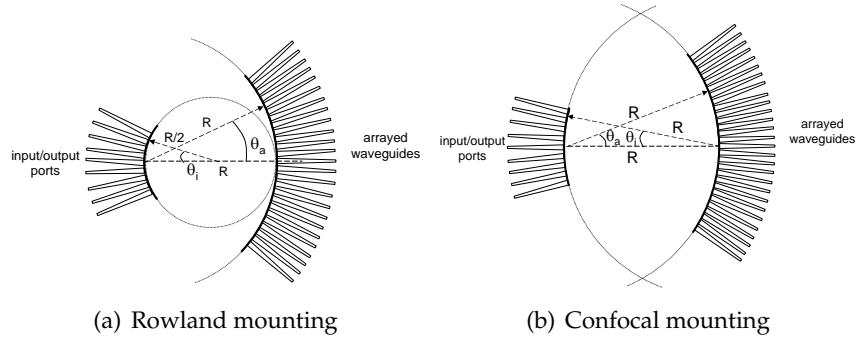
Figure 7.2: Schematics of star coupler mountings

higher order derivatives. This section considers first-order imaging only, considering the grating equation and no aberrations. Usually, the grating and focus lines are dimensioned as circles. Multiple configurations are possible, such as the Rowland mounting, in which the focus line is on a circle half the radius of the grating line, and the grating line is located on the focus line circle. This is illustrated in figure 7.2(a). For a planar grating, the Rowland configuration avoids second order aberration (defocus) and third order aberration (coma) [76]. However, as the first AWG devices originated from connecting optimised star couplers [77], a lot of AWGs are designed with a confocal configuration (figure 7.2(b)). For low port counts, the difference of the input and output port positions is small. However, for larger port counts the confocal configuration is less ideal with respect to aberrations.

Normally, the delay length between arrayed waveguides is a set to an integer multiple of a unit delay length ΔL, which, in combination with the

---

The optical path function can be expanded into a power series:

$$F = \sum_{n=1}^{\infty} \frac{1}{n!} F^{(n)} w^n$$

$F^{(n)}(0)$ are the aberration coefficients. From $F^{(1)}(0) = 0$ (Fermat's principle) we can obtain the grating equation:

$$n_s r_i^{(1)}(0) + n_g L^{(1)}(0) + G^{(1)}(0) m' \lambda + n_s r_j^{(1)}(0) = 0 \tag{7.2}$$

$F^{(2)}(0)$, $F^{(3)}(0)$ and $F^{(4)}(0)$ correspond to defocus, coma and spherical aberrations, respectively.

propagation length differences in the star couplers, defines the wavelength of maximum transmission between two ports, as

$$m'\lambda + n_g\Delta L - n_f R\Delta\alpha\left(\sin\theta_i + \sin\theta_o\right) = 0 \tag{7.3}$$

with n_g the effective index in the wires of which the delay length exists, n_f the slab index, R the radius of the grating line, $\Delta\alpha$ the angular pitch of the grating arms, m' an integer and θ_i and θ_o the input and output angles, respectively.

For the central input port and output ports ($\theta_i = \theta_o = 0$), the dispersion of the output angle is obtained by differentiation of eq. 7.3. As WDM works with equidistantly spaced frequency channels, rather than equidistant wavelength channels, we will work with the frequency as the main parameter.

$$D_\theta = \frac{d\theta_o}{df} = -\frac{n_{gg}\Delta L}{f n_f R\Delta\alpha} \tag{7.4}$$

In a paraxial approximation, this equation also holds for all diagonal ports ($\theta_i = -\theta_o$). For other paths, a small correction term needs to be added. Still, the dispersion increases for larger ΔL and for smaller arrayed waveguide pitches[2].

The angular dispersion and the frequency channel spacing define the angular location of the input and output ports, with respect to the focal point of the focus line (on the grating line), and the symmetry axis of the star coupler.

The linear dispersion Δ_l on the focus line is given by $R\Delta_\theta$, both for confocal and Rowland star couplers, as for a Rowland mounting, the angle from the centre of curvature of the focus line is equal to $2\theta_o$, and the radius of the circle is $\frac{R}{2}$. So, the linear dispersion is

$$D_l = -\frac{n_{gg}\Delta L}{f n_f \Delta\alpha} \tag{7.5}$$

The relation between the central frequency, which is transmitted between diagonal input and output ports, and the delay length is

$$f_0 = \frac{c}{n_g\Delta L}m, \tag{7.6}$$

with $m = -m'$ the grating order.

---

[2]Note that this can easily be understood from the nature of the dispersive phenomenon. In the Fraunhofer diffraction approximation, the Fourier transform converts the linear phase at the aperture side into a displacement in the far field. A stronger phase deviation will thus be converted into a larger displacement in the far field

The frequency free spectral range of the device is, again for all diagonal output ports,

$$\Delta f_{FSR} = \frac{c}{n_{gg}\Delta L}. \tag{7.7}$$

Note the use of the group index instead of the effective index. Whether eq. 7.6 or eq. 7.7 is used to set the delay length depends on the type of device. If ΔL is set from the desired FSR, the central frequency can only be an integer multiple of the FSR, corrected for the waveguide dispersion. In the other case, the FSR can only be an integer fraction of the central frequency, again corrected for waveguide dispersion. For non-diagonal input ports, the FSR is slightly different.

As the field at the input side is periodically sampled, the resulting field at the output focal line will be periodical too. In other words, like every grating, the AWG has multiple diffraction orders. In a paraxial approximation, the angular period (spatial FSR) of the output field is given by

$$\Delta\theta_{SFSR} \equiv \frac{\lambda}{n_f R \Delta\alpha} = |D_\theta \Delta f_{FSR}|. \tag{7.8}$$

The far field in the output star coupler, just in front of the output waveguides, is approximately given by the product of an array factor, expressing this periodicity, and an envelope factor, which is the far field of an individual array waveguide aperture. The array factor can be expressed as

$$U_a = \sum_m U_i (md_a) e^{j2\pi m \frac{\theta}{\Delta\theta_{SFSR}}}, \tag{7.9}$$

with $U_i (md_a)$ the sampled far field at the input of the array and m the grating order.

The width of the lobes in the array factor is influenced by the number of waveguides. A higher number of waveguides will lead to smaller lobes.

The passband width is influenced by the width of the output waveguide aperture relative to the output waveguide pitch. In a Gaussian approximation for the modal fields, and within the Fraunhofer approximation, the L-dB bandwidth is approximately given by [78]

$$\Delta f_{L,gaussian} \approx 0.77 \frac{w_e}{D}\sqrt{L}, \tag{7.10}$$

with w_e the effective Gaussian mode width.

7.1.1 Loss

In an AWG, there are both intrinsic and fabrication-defined sources of loss. As we will see, in the next section, in a high index contrast AWG, these losses are higher and care should be taken to keep them low.

Truncation loss In the input star coupler, part of the diffracted light from an input aperture is not captured by the finite array aperture, resulting in truncation loss. The lost part contains the high spatial frequencies of the field at the input port aperture, and has an exponentially decreasing nature for typical waveguides. This loss can be limited to great extent by choosing a wide enough array aperture, or by tailoring the aperture field profile.

Loss non-uniformity The envelope factor containing the information about a single aperture diminishes away from the output port to which the central frequency is coupled, which results in insertion loss non-uniformity between the frequency channels. Keeping this non-uniformity low requires a slowly varying envelope factor within the spatial free spectral range. The N_{ch} channels need to fill one spatial period. As at the edge of the spatial period, relative to the central channel, the two diffraction orders must carry approximately the same power, the outer channels have an excess loss of about 3 dB.

Coupling to adjacent diffraction orders As part of the light is inevitably coupled to adjacent diffraction orders, this effect mainly determines the minimum insertion loss (the loss of the central channel). Therefore, the adjacent orders should be excited as less as possible, by properly choosing the array and envelope factor in the far field pattern, within the limits of choice. One option is to increase the number of arrayed waveguides in order to decrease the width of the lobes in the array factor. Alternatively, the arrayed waveguide aperture width can be increased in order to decrease the width of the envelope factor. However, this also increases the loss non-uniformity and decreases the spatial periodicity of the far field.

Reflection and scattering loss As the lithographical resolution is finite, small gaps will be present between the arrayed waveguide apertures. Light scatters from the interface between the slab and these gaps, giving rise to loss. Therefore, the gaps should be kept as small as possible within the resolution limit of the fabrication process.

7.1.2 Crosstalk

Intrinsic crosstalk is generated in all AWG devices by truncation and modal coupling, which for communication applications will put the major limits on the design parameters. A design will typically start from the crosstalk requirement.

Truncation Together with power loss due to truncation at the finite array aperture, part of the information about the wave at the input aperture is lost, resulting in imperfect imaging and thus crosstalk into the neighbouring channels at the output side[3]. In a worst-case scenario, the direct neighbours of the intended output port each receive half of this crosstalk power, which is equal to the lost power. In this way, a 0.01 dB truncation loss will maximally be joined by a -29.4 dB nearest neighbour crosstalk, 0.001 dB power loss with -39.4 dB crosstalk and 0.0001 dB with -49.4 dB crosstalk.

Coupling at the port apertures The overlap between the imaged field and the mode tails of the output ports neighbouring the intended output port can give rise to crosstalk. If the AWG is designed symmetrically, it can be seen from reciprocity this crosstalk does not occur, as an equal coupling will happen at the input side. However, this holds only with perfect imaging. With imperfect imaging, crosstalk still results from coupling between ports. Thus, the linear spacing between ports should be sufficiently large. This puts a requirement on the angular dispersion or the radius of the focal line.

Coupling of the arrayed waveguides As the arrayed waveguides are coupled near the star couplers, there is some transfer of power between them, both at the beginning and end points of the array. This distorts the power distribution among the waveguides and thus introduces crosstalk at the output ports.

Coupling to adjacent orders In order to avoid high crosstalk due to coupling of the adjacent orders to output waveguides, the arrayed waveguide pitch should thus be small enough, and the total extent of the output

---

[3]An equivalent way of viewing this is to regard refer the truncation of the field to the output side of the array, and seeing this as the reciprocal effect of the truncation at the input side.

waveguides should not exceed the spatial FSR. As the envelope factor defines the relative strength of the adjacent orders, the envelope should be high within a single spatial FSR and as low as possible where overlapping between the adjacent orders.

Phase errors   If the effective index of the waveguides is altered in a deterministic or stochastic manner by waveguide width or height variations or by refractive index variations, phase errors are introduced in each arm, leading to a difference from the intended phase variation at the grating line, and thus crosstalk. Generally, this is the limiting contribution to crosstalk for most AWG devices.

7.1.3   Tapering of the port and array apertures

For several reasons, the narrow array wires should be tapered towards the slab of the star coupler. At the array entrance and exit, scattering occurs a the gaps between the waveguides. Therefore, the gaps should be as small as possible, coupling the arrayed waveguides. Secondly, the envelope factor should be designed so that the power radiated to lost diffraction orders is minimal [79]. Due to reciprocity, this holds not only at the array exit side but also at the entrance side. For simple waveguide modes, this will however also increase the loss non-uniformity. Tapering of the port waveguides is not always done, and depends on the bandwidth and crosstalk requirements, defining the port spacing and its fill factor.

7.1.4   Illumination of the array and passband flattening

The uniformity of the array excitation has a great impact on crosstalk, phase transfer and flatness of the transmission band. With a flat, rectangular-like illumination, the truncation loss and crosstalk can be lowered. However, if the excitation of the array is uniform, then the imaged field, being approximately the Fourier transform of the array aperture field, will be sinc-like, exhibiting larger sidelobes and a smaller linear phase response bandwidth than with a less uniform excitation. Inversely, a sinc-like field at the array exit aperture will give a more flat passband. Such a field can be achieved by inserting phase delays in the arms, or by exciting the array as such.

   Some options for passband flattening by shaping the beam at the input side have been explored, such as obtaining a slightly double-lobed aperture field through small MMIs [80] or parabolic horns [81]. At the receiver side, the waveguides were normal waveguides. A drawback of this

approach may be the passband ripple, and the asymmetry of the configuration, although reciprocity still holds of course.

With a high index contrast system, beam shaping is possible with very short tapers [82], opening up options for beamshaping at the input (and output) side of the AWG.

Other options for passband flattening include altering the amplitude or phase in the array waveguides as already mentioned, adding interleavers with a very rectangular spectrum in front of multiple AWG devices, or, for demultiplexers, by adding a Mach-Zehnder Interferometer (MZI) in front of the AWG [83].

7.1.5   Local contrast reduction

In order to reduce both reflection and scattering losses and losses to the other diffraction orders, the index contrast can be locally reduced around the star coupler. This can for instance be done by using a more shallow etch and converting between the deep and shallow etch adiabatically, through a double etch scheme [84]. We will point out further on how this was done in the SOI system.

7.1.6   Polarisation dependence

Polarisation dependence results from polarisation dependent waveguides and possibly from a polarisation dependent slab waveguide. The effect is a wavelength shift, which is detrimental for communication applications, and additional loss if the modes experience a different propagation loss. Birefringence even exists in low-index-contrast waveguides, but the problem is generally much more severe with higher index contrasts. A number of methods exist to make an AWG polarisation independent, the most obvious being to make the waveguides non-birefringent. For demultiplexer devices with higher index contrast waveguides, a generally more achievable solution is polarisation diversity. In this scheme, the polarisations are split at the entrance, and the shift between the input port waveguides is chosen equal to the polarisation dependent shift on the focus line, making the two inputs arrive in the same receiver waveguide. Another option is to compensate for the birefringence by insertion of a section with different birefringence in each waveguide of the array. This results for instance in a triangular overlayer, as the length of the compensation section should grow linearly with the waveguide number. A fourth possibility is to let the diffraction order for one polarisation overlap with a neighbouring diffraction order of the other polarisation, in this way, the diffraction

orders for both polarisations are different, but still the AWG is polarisation independent. However, this limits the choice of the achievable FSR and delay lengths. A last method is to insert a polarisation rotator into the middle of the array, converting between both polarisations. However, a good polarisation rotating device is needed of course, which is not evident in semiconductor waveguides.

7.2 General design procedure

For use as a router or demultiplexer for communication purposes with WDM channel spacings, one usually starts from a crosstalk requirement. However, the design details are different for a router and a demultiplexer. If the number of input N_{in} or output ports N_{out} is even, then an imaginary central port is used during the design for the relevant star coupler.

7.2.1 Wavelength router

For use as a wavelength routing device with N_{ch} channels, which was the primary focus of the work set out in chapter 9, a cyclic AWG is desired, with $N_{in} = N_{out} = N_{ch}$. In such an AWG, the frequency channels connecting input and output ports wrap around as indicated in figure 7.3. Therefore, the first step is to set the FSR as

$$\Delta f_{FSR} = N_{ch} \cdot \Delta f_{ch}, \tag{7.11}$$

with Δf_{ch} the channel spacing. This defines the delay length through eq. 7.7.

A possible design procedure is given in figure 7.4. The minimum gap width and the minimum waveguide aperture width define the minimum linear array pitch Δx_a. For the ports, a minimum linear pitch $\Delta x_{p,min}$ and aperture width is set to obtain a certain bandwidth and simultaneously avoid crosstalk due to coupling. The choice of the aperture width then defines the angular array aperture through the diffraction angle and the requirement on loss and crosstalk due to truncation. For instance, a 4σ angle[4] can be chosen. The choice of port pitch, together with the channel spacing, defines the required linear dispersion. From the linear dispersion and the already fixed ΔL, the angular array waveguide pitch is

---

[4]Here, the σ angle is defined as the two-sided Gaussian equivalent angle, for which the intensity is reduced at both sides with a factor $\mathrm{e}^{-\frac{1^2}{2}}$ relative to the maximum far field intensity, 2σ as $\mathrm{e}^{-\frac{2^2}{2}}$ and so on.

Figure 7.3: Transmission scheme of AWG router with three input and output ports
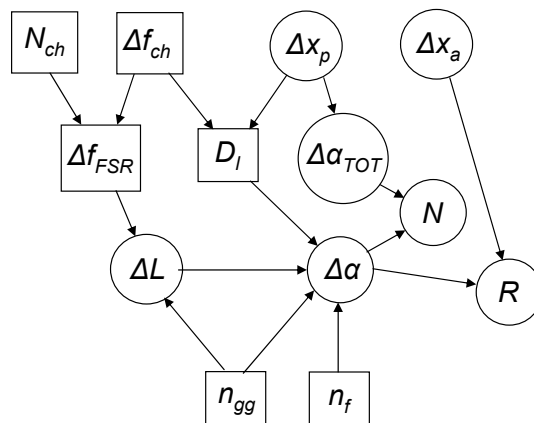
Figure 7.4: Possible design flow for a cyclic AWG router.

fixed through eq. 7.5, and the number of arrayed waveguides N is known. From the minimum linear array pitch, and the number of waveguides, the radius of the grating line is set. It can be chosen larger to use broader waveguide apertures.

Combining the formulas, the radius can directly be written as

$$R = N_{ch} \frac{n_f}{\lambda} \Delta x_a \Delta x_{p,min}. \tag{7.12}$$

Several other options are possible. For instance, the number of arrayed waveguides can be set after the minimum gap width as well as port and arrayed waveguide aperture widths are chosen. In this way, the width of the array factor lobes can be adjusted. Given the total angular array aperture, known from the port aperture width, this sets both $\Delta\alpha$ and R. The angular dispersion is then known, and the port angles can be calculated and checked to see if the resulting port spacings are at least as large as the minimum spacing set.

For a wavelength router, the FSR and SFSR exactly encompass the N_{ch} channels needed. Therefore, the maximum loss non-uniformity will inevitably be about 3 dB.

In conclusion, for a router we start from the cyclic and crosstalk requirements and all free physical parameters are derived from these.

7.2.2 Demultiplexer

When the AWG is designed as a demultiplexer, we start from the loss non-uniformity and crosstalk requirements. In order to achieve a low non-uniformity, the FSR is taken much larger than the actually $N_{ch}\Delta f_{ch}$ needed. The rest of the design procedure is equal to that of a router.

7.3   AWG devices in SOI

The published work on AWG devices in SOI photonic wires is still very limited. T. Fukazawa, F. Ohno and co-workers investigated extremely compact devices with large channel spacings of 6 to 11 nm and a FSR larger than 90 nm. Their first AWG [85] still had large crosstalk of -5 dB and large loss of 12 dB. Later, an extremely compact AWG was demonstrated [86] with a footprint of 70 μm ×60 μm, with reduced insertion losses due to optimized tapers, but still with a crosstalk of -8 dB to -13 dB. In a third report [87], the authors showed a crosstalk reduction, however through a reduction of the filter order (larger FSR). Simulation work has been done by D. Dai and S. He [88] on compact AWG form factors, as well as fabrication in SOI with amorphous silicon [89].

7.3.1   Form factors

In low and medium index contrast material systems, the AWG size is mainly limited by the bend radius. Most AWG devices use a variant of the layout used by Smit [78], Dragone [77] and others, although Takahashi et al.[90] for instance used the box-like layout used in this work, for de-multiplexing 10GHz spaced channels with a GeO_2 - SiO_2/SiO_2 waveguide system.

With a high index contrast, other design options are possible as the bend radius can be greatly reduced. The possibilities are limited by the delay length needed and the waveguide spacing needed for decoupling the waveguides. This spacing is discussed in section 2.4.5. The most conventional option is to use a rectangular layout, that was also used in low-index-contrast waveguides. It allows for an elegant design with a close packing. The minimum delay length is equal to three times the isolated waveguide spacing.

This is the configuration used throughout this work. However, with the very small bend radius offered by SOI, other packing options become possible, by *folding* the delay lengths. For instance, each waveguide can follow a snake-like pattern, allowing for a very close packing [88]. However, the approach proposed by the authors in [88] suffers some non-trivial problems. The total length is corrected for the difference in refractive and group index, but the number of bends does not scale linearly, which makes it difficult to achieve both the desired total group delay and phase delay. A general problem with folding delay lines is that phase errors become more likely, due to mask digitisation and due to the difference in effective indices between bent and straight wires. Still, extremely compact AWG

devices are possible, so for limited performance requirements, folding delay lines may be a good idea.

7.3.2   Index contrast reduction

To locally reduce the index contrast around the star couplers in SOI, a double etch scheme was used. Near the apertures, the waveguides are etched only 70 nm deep. While this drastically reduces the index contrast, note that the resulting contrast is still high compared to low-index-contrast material systems. Both the shallow and deeply etched waveguides are tapered towards the slab waveguide. This is illustrated in figure 7.5. The grating couplers for interfacing to fibre are etched 70 nm deep as well, allowing both structures to be on the same mask and be fabricated in a single step. The alignment tolerances of the stepper (better than 50 nm) are good enough so that no excess losses are introduced.

The shallow etch should be used at the array side. At the input and output port sides, a shallow etch is more of a downside, unless it would reduce interface losses. This can be seen as follows. A shallow etch requires a larger distance between the ports to avoid coupling. Therefore, the required linear dispersion is higher and thus the angular pitch of the array should be smaller, which leads to a larger radius as the linear array pitch is defined by the chosen aperture width and gap. With a smaller diffraction angle but a smaller angular pitch, there is no gain on the number of waveguides. The bandwidth, defined by the fill factor of the port pitch, will be about the same regardless shallow or deep etch. In conclusion, with a larger radius, and room needed for tapering between deep and shallow etch, a shallow port etch only increases the AWG footprint and does not influence performance unless there is a limited improvement in losses.

7.3.3   Main design and physical parameters

Free spectral range

With a minimum waveguide spacing for isolation of 1 µm, the minimum delay length is 3 µm, and the maximum achievable FSR is 22 THz or 178 nm in a constant group index approximation. Over such a long range however, the second order dispersion will be non-negligible and the channel spacing will deviate from the designed value.

Figure 7.5: Schematic of a local shallow etch for index contrast reduction in the star coupler access waveguides.

(a) Ports side · (b) Array side

Figure 7.6: SEM pictures of a local shallow etch for index contrast reduction in the star coupler access waveguides.

(a) deep tapers   (b) shallow tapers

Figure 7.7: Simulated overlap of isolated aperture modal fields for oxide-clad star couplers. Discontinuities are due to the limited resolution of the width parameter.

Aperture width and spacing

In a high vertical index contrast system such as SOI, an additional reason for tapering both array and port waveguides is to avoid vertical radiation loss and reflections due to vertical mismatch, as noted further on. Still, the main reason is to adjust the passband width.

The array aperture width can be taken 1 to 1.5 μm. With a minimum controllable gap of around 180 nm, the pitch is thus 1.18 to 1.68 nm, corresponding to an angular SFSR of 18.7°to 26.6°.

Figure 7.7 shows the overlap of the aperture modes for both deep and shallow etched star couplers, as a function of the aperture width and the gap. For deep port tapers, a safe overlap is obtained with a gap of 800 nm for any aperture width. The overlap decreases with waveguide width, so the filling ratio of the port spacing can be chosen large to obtain a large bandwidth. This also holds for shallow waveguides, but the gap should at least be 2 μm.

Focus distance

Designing the focus distance through eq. 7.12, typical values of the radius are given in table 7.1. For a large number of channels, the focus distance becomes a limiting factor to the AWG footprint. With a shallow etch, the diffraction angle becomes smaller and the radius larger for the same number of array waveguides.

N_{ch}	4	8	16	32	64
$\Delta x_p = 1.8\,\mu m$	19.7 µm	39.4 µm	78.9 µm	158 µm	316 µm
$\Delta x_p = 2.5\,\mu m$	27.4 µm	54.8 µm	110 µm	219 µm	438 µm

Table 7.1: Star coupler focus distance for an AWG designed as a cyclic router, as a function of channel number, for small and wide port spacing. The arrayed waveguide spacing is 1.5 µm.

7.3.4 Sources of loss and crosstalk

In a high index contrast system, basically all sources of loss and crosstalk are more severe than in a low-index-contrast system. The narrow waveguides are tapered open towards the slab waveguide of the star coupler. Nevertheless, the diffraction angle is larger than for lower index contrast systems. Therefore, the star coupler radius will be smaller or the number of waveguides larger in order to keep the loss and crosstalk due to truncation under control.

High index contrast waveguides have a strong discontinuity in the electric field normal to the waveguide core sidewall. This gives rise to high spatial frequency components, that relate to tails of the diffracted far field that are larger than for lower index contrast systems. Therefore, high index contrast AWGs, and SOI AWGs in particular will have a larger crosstalk floor due to truncation.

Phase errors in waveguides are stronger in higher index contrast devices, as we have already seen in chapter 6 with the Mach-Zehdner lattice filters. Therefore, an extreme dimensional control is needed, or tuning of every waveguide. In an AWG with WDM channel spacings, the longest waveguides are of the order of a mm long, giving rise to a high phase error. Even for the waveguides carrying the bulk of the power, the phase error can become large and thus a large crosstalk results. One possibility is of course to use a larger FSR, for routers with large CWDM channel spacing [87]. However, DWDM intrinsically suffers phase errors as a crosstalk-limiting factor.

7.3.5 The vertical dimension

Very specific to the high vertical index contrast of SOI are the vertical mode mismatches that appear in a component at all junctions between structures of different width (see chapter 2, figure 2.7). Due to this vertical mismatch, light scatters not only to the intended guide mode, but also in reflection

to the counterpropagating mode, to horizontally or vertically higher order guided modes if present or to radiative modes. So, every non-adiabatic width change generates loss and reflections, which are not completely negligible as in lower vertical index contrast systems. In AWG devices, this happens at the apertures with the slab. However, as the apertures are tapered to 1 to 2 μm, the vertical mismatch is very small to negligible. Nevertheless, as we will see, an effective index approximation still does not predict the mode tails adequately. As these are important with respect to crosstalk and loss, the real 3-D mode profiles should be used. Also, if coupling between the waveguides were taken into account, 3-D simulations are necessary in order to predict the influence. This is a labour- and computation- intensive job.

7.3.6 The slab waveguide an Si layer thickness influence

As seen in section 2.4.1, the slab waveguide mode has an effective index of around 2.82, but a group index of 3.7. While a lot of reports on AWG theory ignore the dispersion of the slab, it is therefore useful to check its influence. A quick calculation shows however that the difference in dispersion of the AWG is smaller than 10^{-4} degrees/GHz, even for transmission between the outer ports.

Secondly, Si layer thickness variations can have an influence. Stochastic roughness will both give rise to small phase noise in the slab, which can be expected to be very small however, and a mean slab thickness that is different from the design value. The silicon layer thickness can also deviate from the designed value both wafer-to-wafer and intra-wafer (see chapter 3). This influences both the slab refractive index and the effective and group index of the array waveguides. Differentiating eq. 7.4 to the height h, the relative deviation of the dispersion is given by

$$\frac{\frac{D_\theta}{dh}}{D_\theta} = \frac{1}{n_{gg}}\frac{dn_{gg}}{dh} - \frac{1}{n_f}\frac{dn_f}{dh}. \tag{7.13}$$

With $\frac{dn_{gg}}{dh} \approx 1.6 \cdot 10^{-3}/nm$ and $\frac{dn_f}{dh} \approx 3.3 \cdot 10^{-3}/nm$, both effects on the AWG dispersion are in opposite direction, and therefore the net effect is limited to $\frac{\frac{D_\theta}{dh}}{D_\theta} \approx 5 \cdot 10^{-4}/nm$.

With a typical output port pitch of a few micrometre, mean height differences of the order of 1 nm due to stochastic deviations, would induce a deviation of the optimum port location of only about 1 nm. Likewise, wafer-to-wafer thickness variations of 10 to 20 nm would give a difference

in optimum port location of about 10 nm, which will have a probably noticeable but nevertheless small effect.

Designing the focus distance through eq. 7.12, the relative defocus is of the order of $10^{-3}/nm$, in absolute values some tens to hundreds of nm depending on the number of channels.

For the design of the AWG devices in this work, the dispersion of the slab was taken into account, by adapting the port angles in order to optimize transmission when exciting the central input port.

7.3.7 Polarisation dependence

The polarisation dependence of SOI AWG devices is extremely large. With the waveguides used here, the polarisation dependent shift is of the order of 300 nm. The difference in order between the TE-like and TM-like modes for the same wavelength is some tens to hundreds. Moreover, for narrow wires, their group index is drastically different too, leading to a different free spectral range. Therefore, the techniques developed for low-index-contrast for obtaining a polarisation independent device fail.

For oxide clad wires around $\lambda = 1550\,nm$, at a width of 690 nm, the group indices of the TE-like and TM-like mode are identical. So, if the delay length is implemented in such wires, both modes have the same FSR, and tailoring the waveguides so that the diffraction orders overlap is possible. At a width of 690 nm, the waveguide is multimode, so the bend sections should still be kept narrow, with tapers converting between the broad and narrow waveguides. However, the wire width of equal group index is very intolerant to waveguide height variations and to wavelength variations. With an air top cladding, this solution does not work, as around 700 nm the TM0 and TE1 mode interact and become very hybrid.

Therefore, there are only two options left for obtaining polarisation independent functioning. One is to use square-core, oxide clad wires. Although the polarisation independent behaviour is very intolerant to small dimensional variations, it should be possible to compensate for this birefringence using one of the methods outlaid above. Still, very tight control will be needed, and adaptable from sample to sample. The second possibility is to use polarisation diversity and converting the TM-like mode to a TE-like mode, before entering a second identical structure or the same structure. This can be done through 2D fibre couplers as explained in section 2.4.7. An example of a polarisation independent operating AWG based device is given in chapter 9, where the two paths corresponding with the orthogonal polarisations in the fibre are looped in opposite direction through the same AWG.

7.4   Simulation method

Before discussing some simulation results, the simulation method used is presented. Basically, several numerical methods were grouped together and implemented in a software package for simulation of SOI arrayed waveguide gratings.

7.4.1   The complete picture

It is currently impossible to fully simulate even a small AWG device with a numerical method such as FDTD due to the huge amount of memory and time needed, and this situation will persist in the nearby future. Therefore, the device is broken into pieces that are individually simulated. The AWG is very suited for this, with two star couplers and a waveguide array. For $N \times N$ router devices, the star couplers are even identical. Transmission matrices can be obtained for the different parts, and the full transmission is obtained by simple multiplication.

   Simulation of the uncoupled arrayed waveguides is straightforward. If all bend sections are the same in each waveguide and the delay is only composed of straight waveguides, the wavelength-dependent effective index and loss give enough information and simulation simply consists of multiplying the mode with the right phase and loss factors. In the case of the delay lines containing bends, the situation gets somewhat more complicated but still, the mode is just multiplied with a series of phase factors.

7.4.2   Star coupler transmission

Simulation of the star couplers can be calculated with different levels of approximation and accuracy. For low-index-contrast waveguides, the waveguide modes can be approximated by Gaussians. In the Fraunhofer approximation, the far field is obtained by a Fourier transform. In the case of Gaussian mode profiles, the field at the array aperture is thus a Gaussian too, and an overlap can be taken with the waveguide modes. However, for high-index contrast waveguides, the Gaussian approximation seriously underestimates the energy in the mode tails, and due to the wide diffraction angle of narrow waveguides, the paraxial approximation is not valid anymore except for the inner waveguides. Simply calculating the far field as the Fourier transform of the field in the aperture is too simple. Therefore, the two-dimensional procedure outlined in [91] was used in order to calculate transmission between a port and array waveguide. First, the isolated waveguide modes are propagated to the apertures at both the

Figure 7.8: Definitions used for calculating the transmission between port and array waveguides in a star coupler.

port and array side. With a high vertical index contrast, this simulation should be done in three dimensions to reliably simulate the effect of possible coupling (see further).

For the transmission between the apertures, the definitions are given in figure 7.8, with in the example $\theta_i < 0$, $\alpha_k > 0$, $\theta_A > 0$ and $\theta_B > 0$. First, the waveguide profiles are corrected in phase for their rotation with respect to the propagation vector $\boldsymbol{AB}$ between the waveguides. At the input side, the phase-corrected mode is expanded in in-plane plane waves through a Fourier transform. These plane waves are are propagated through the slab, and added back together at the opposite side of the star coupler through an inverse Fourier transform. There, the overlap with the phase-corrected mode of the local array aperture is taken in order to arrive at the transmission from waveguide to waveguide. No paraxial (Fresnel) or far-field (Fraunhofer) approximations are made. With the definitions of figure 7.8, $\theta_A = -\theta_i$ and $\theta_B = \alpha_k$, and the correction angles are $\theta_A - \angle AB$ and $-\theta_B + \angle AB$. The propagation vector $\boldsymbol{AB}$ can be calculated easily for both confocal and Rowland type mountings.

In these methods, all propagation is considered to be in-plane. However, the real three-dimensional field expansion will contain terms (plane waves) with a partly vertical propagation direction. This fact, and the influence of the vertical structure on this, are neglected. As it is currently impossible to simulate a star coupler with FDTD, it is impossible to check the error made by the two-dimensional (or 2.5D) approach used.

A second note is that the field profiles just in front of the aperture are used, and not the near-field just after the aperture, to calculate the far field. This is of course another approximation, which gets better with waveguide width.

Finally, both amplitude and phase information of the star coupler transmission should be taken into account for accurate simulation.

7.4.3   Waveguide modes

Often, the waveguide modes and fields are approximated as Gaussians, allowing simple reasoning and approximate numerical models. For a low-index-contrast system, this Gaussian approximation is good. However, as the index contrast rises, the tails of the Gaussian field become a less good approximation for the real exponential mode tails. An accurate representation of these tails is needed in order to simulate truncation and waveguide coupling. Additionally, in a photonic wire, there is a strong discontinuity in the electric field perpendicular to the sidewall. In the star coupler, this discontinuity will have its influence in the far field.

In the arrayed waveguides, propagation is sufficiently represented by the effective index as a function of wavelength, if all waveguides contain the same bends. Near the star couplers, an actual field representation is needed. An interesting approach is to approximate the waveguide modes by a sum of several Gaussians [92]. Such a sum may be more representative for the mode tails, but still allows for easy reasoning and calculations. However, for accurate simulation, computer force today allows for partly using actual the actual mode profiles.

2.5-D approximation   For a low vertical index-contrast material system, the system can be reduced to two dimensions through an effective index approximation. The resulting one-dimensional mode profile can then be used in the calculation. What profile to use with a high vertical index contrast ? One option is to use the effective index approximation too, which is still fairly acceptable for broad isolated apertures. However, with coupled apertures, the effective index method is too approximate. In order to reduce the 2D field profile resulting from a full 3D simulation of the possibly coupled tapers, several methods are possible. For deeply etched apertures, definitely in the symmetric case with an oxide upper cladding, it is possible to use a slice of the 2D field in the horizontal plane, through the middle of the wire. Of course this method is not very rigorous, and it does not work for shallowly etched tapers. Therefore, a possibly better approach is the following [93]: both at the port and array side, the field is

sliced in vertical slices. For each vertical slice of the field the overlap with the slab mode can be calculated. Each number on the horizontal axis thus describes the vertical behaviour at that point and a 1D representation of the 2D field profile is obtained. This 1D representation can then be used in the calculation outlined above.

7.4.4 Coupling of the arrayed waveguides

The total array aperture should approximate the slab waveguide as well as possible in order to avoid insertion loss. However, this intrinsically means the arrayed waveguides are coupled near the star couplers, and one cannot simply take the overlap of the far field with the isolated aperture modes to accurately simulate the star coupler transmission. For the calculations in this chapter, the coupling was not properly taken into account, leading to a too low calculated insertion loss and a deviation of the insertion loss non-uniformity.

A better simulation approach is the method laid out by Klekamp and Münzner [94], explained in the next paragraphs, but was not yet used for the calculations shown further on. In an approximation, one excites a single waveguide with the isolated waveguide mode field, starting from a position where the waveguides are reasonably well uncoupled, and propagate to the star coupler boundary. The resulting field, the Green's function of the array [94], can then be used as an expansion basis for the far field:

$$U_{ff}(x) = \sum_{i=-N-N_c}^{N+N_c} \alpha_i G_i(x - x_i),$$ (7.14)

where the Green's function extends over N_c waveguides at both sides of the originally excited waveguide. For a constant $\Delta\alpha$, constant waveguide profiles and with dummy waveguides at the outer sides, $G_i = G$.

The Green's function can in its turn be expanded into the modes of the isolated apertures in order to obtain the star coupler transmission. However, the isolated waveguide aperture modes do not form an orthonormal basis for the total array aperture, as the waveguides are coupled. This can be solved by taking the overlap matrix of the isolated modes into account. Klekamp and Münzner have proven [94] that, given the isolated mode fields Φ_i of the array apertures, and given the scalar product (power overlap)

$$(\Phi_i|\Phi_j) := \int \overline{\Phi_i(\rho)}\Phi_j(\rho)\,d\rho,$$ (7.15)

the expansion of an element F of the vector space spanned by the array modes Φ

$$F = \sum_i \eta_i \Phi_i, \tag{7.16}$$

is defined by the expansion coefficients

$$\eta = (\Phi|F)\left(A^T\right)^{-1}, \tag{7.17}$$

with $A_{ij} = (\Phi_i|\Phi_j)$. With constant $\Delta\alpha$ and waveguide profiles, the elements of A depend on the distance between the waveguides only.

However, when actually implementing this method, some numerical problems arise, making it not straightforward to obtain reliable simulation results.

7.4.5   Implementation

The method for calculation of the star coupler and AWG transmission discussed above was implemented in a computer programme. The mode fields were calculated using the mode expansion software CAMFR or Fimmwave, and then further manipulated [5]. Either the effective index approximation was used, or the 3-D mode fields were used. The 3-D mode profiles were then either used through the method described in [93] or by taking a horizontal slice through the profile. The latter method proves to be a good approximation for both deep and shallowly etched waveguide apertures, while reducing the computation time. Coupling between waveguides was not simulated or corrected for. This will appear in the following results as an unrealistic insertion loss for narrow apertures, and an unrealistic crosstalk floor if other crosstalk generating factors are smaller. Accurate material models were used in order to accurately incorporate the wavelength dependence of the waveguides and structure. For all simulations, the calculated modes were extended and used a 2^{15} points (32768) and a $\frac{\lambda}{n_f 40}$ spacing, in order to both obtain a high resolution and a sufficiently length of the far field, respectively. In order to simulate the influence of phase errors in the arrayed waveguides, the phase error model from appendix A, also used in chapter 6, was implemented. The computer code accepts a standard deviation σ_{dw} and a correlation length L_c of the waveguide width variations and calculates the phase error based on the dependency of effective index on waveguide width as reported in chapter 2.

---

[5]Generally, the python scripting language was used, with the help of numerical and scientific libraries for matrix algebra, interpolation, integration and plotting. The used tools were python (www.python.org), NumPy/SciPy (www.scipy.org) and Matplotlib (matplotlib.sourceforge.org)

7.5   Simulation results

In this section, we will discuss the transmission of star couplers, the insertion loss of AWG routing devices and the crosstalk, with and without phase errors in the arrayed waveguides. The simulations shown are limited, but in the next section, experimental results are compared with simulations.

7.5.1   Far field and star couplers

In order to test what approximation of the aperture modes to use for the calculation of the star coupler transmission, the far field of a waveguide aperture was calculated using the simulation programme. Figure 7.9 plots the far field intensity of some waveguide apertures, on a plane trough the focus of a star coupler of a realistic AWG device[6]. These plots show two important features. First, the far field tails clearly deviate from a gaussian behaviour. This higher power in the tails is reminiscent from the high spatial frequency components in the near field, notably due to the discontinuity in the x-component of the electric field at the waveguide sidewalls. The high spatial frequency ripple on the simulated transmission in the tails is an artefact due to truncation of the field in the computation. Windowing before the fast Fourier transform can improve this. Secondly, the effective index method can give different mode tails compared to the methods based on the calculated real mode profiles. For narrow apertures, the main lobe is not well represented as well. For shallowly etched ports, the 2D approximation is better (not shown). For all simulations further on, a horizontal slice through the centre of the 3D mode was used.

Figure 7.10 shows the simulated transmission of star couplers of 5-port AWGs with varying aperture widths, for deeply etched apertures. The angular spread of the array was 5σ[7] in order to calculate transmission for a large number of waveguides.

In order to see the influence of the aperture widths, figure 7.11 shows different combinations. The transmission is shown for a constant port width w_p while varying the arrayed waveguide aperture width w_a, for a constant w_a while varying w_p, and for simultaneously changing both widths. The gap between arrayed waveguides is kept at 200 nm, while the

---

[6]This was easier to calculate than the projection on a circular arc

[7]As noted earlier, the σ angle is defined as the two-sided Gaussian equivalent angle, for which the intensity is reduced at both sides with a factor $e^{-\frac{1^2}{2}}$ relative to the maximum far field intensity, 2σ as $e^{-\frac{2^2}{2}}$ and so on.

(a) aperture width 1 µm

(b) aperture width 1.2 µm

Figure 7.9: Simulated far field of a port aperture on a plane through the focus of a star coupler, for deeply etched apertures. Dashed line: 2-D effective index. Other lines: 3D Slice and 3D-slab mode overlap

(a) $w_a = w_p$= 1.0 µm, N_a=33, R=16.5 µm.

(b) aperture width 1.5 µm, N_a=32, R=31.1 µm.
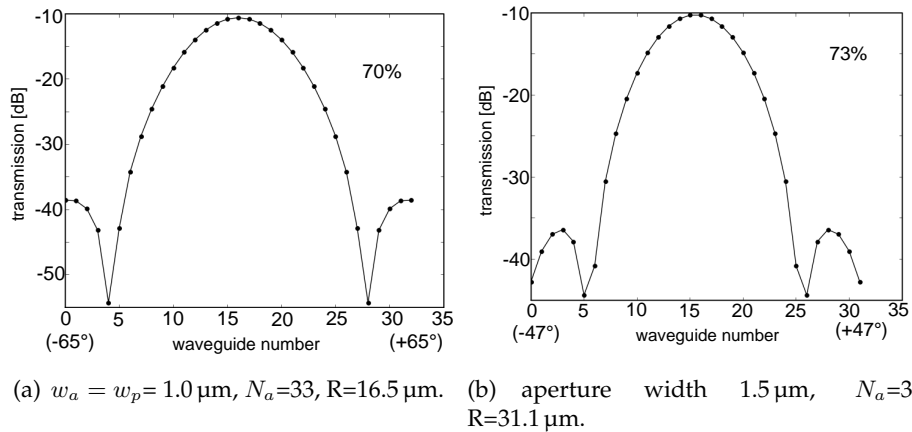
Figure 7.10: Simulated transmission of deeply etched star couplers of AWG devices with 5 ports, for 1.0 and 1.5 µm aperture width. Aperture angle = 5σ, gap between apertures g_a = 200 nm Dashed line: 2-D effective index. Other lines: 3D Slice and 3D-slab mode overlap. The total summed transmission is indicated.

minimum port spacing is 500 nm. Due to the design method, the number of waveguides was approximately kept constant ($\pm$33 for with deeply etched ports, $\pm$25 with shallowly etched ports) and the focus distance is varied to adapt to changing arrayed waveguide aperture width or changing diffraction angle (port width). For the deeply etched port and arrayed waveguide apertures, we see a variation of maximally 0.2 dB as a function of aperture width. For varying port width, the difference is small as the focus is adapted to the diffraction angle. Transmission increases with increasing arrayed waveguide aperture width, although the coupling between the waveguides decreases. This may be explained from the reciprocal situation: with increasing w_a and array pitch, the angular FSR decreases, but at the same time the envelope factor (far field of an isolated aperture) becomes more narrow. In our situation, this leads to an overall decrease of coupling to the nearby diffraction orders and higher coupling to the central diffraction order. From reciprocity, transmission from the input to the arrayed waveguides thus increases. For shallowly etched array apertures however, we see a different behaviour. With a 200 nm gap, the coupling between apertures is much higher than with a deep etch, and therefore the increase of coupling with narrower waveguides dominates the loss behaviour. Simulation for a deep etch with gaps below 50 nm also showed this behaviour to a smaller extent. However, part of the increase in transmission for narrower array apertures is artificial, due to the non-orthonormal expansion basis formed by the isolated modes, as explained above. From figure 7.7, we can see the power overlap between two neighbouring apertures is a few percent, and up to 10% for narrow apertures. Therefore, the transmission and the transmission differences are both overestimated, resulting in transmission greater than unity for narrow apertures. Therefore, only for wide shallowly etched apertures of 1.5 µm and above, transmission will be adequately approximated.

From the measurements of deeply etched star couplers discussed in the next section, it seems we can relatively well simulate the deep star coupler transmission. The decaying tails of the far field, that are truncated at the array aperture are well represented for wider port apertures, and thus an accurate estimate of the truncation loss is obtained. For narrower apertures, the tails might be underestimated.

7.5.2 Loss

As mentioned already, the coupling of the arrayed waveguides near the apertures is not taken into account or corrected for. As the isolated modes do not form an orthonormal basis for representing the field at the array

(a) deep etch

(b) deep port etch, shallow array aperture etch

(c) shallow etch

Figure 7.11: Simulated summed transmission of star couplers of AWG devices with 5 ports, for varying aperture widths. Total array angular spread = 5σ. Simulation for constant port width w_p while varying arrayed waveguide aperture width w_a, for constant w_a while varying w_p and for $w_p = w_a$. The varying width is indicated on the x axis.

interface, the star coupler transmission obtained is overestimated. The overestimation is stronger for stronger coupled apertures. Deeply etched apertures are relatively weakly coupled even for 200 nm gaps. Therefore, the simulated insertion loss will only to a small extent overestimate the actual loss.

In order to obtain the minimum angular spread of the array needed to avoid truncation loss, 9-channel AWG router devices were simulated with varying angular spread, while keeping the focus distance and all other parameters except the number of waveguides constant. Nine ports were chosen to obtain a large enough resolution (number of waveguides per radian) but keep computation time under control. The AWGs were calculated only on the channel wavelengths, and the loss plotted is that between central input and output. As a design for AWG routers was done, the centre frequency was not a design parameter due to the high index contrast, as was explained in chapter 4. Therefore, the real passband wavelength was calculated first. Figure 7.12(a) plots the insertion loss for the test AWG with 200 GHz spacing, for the three combinations of deep and shallow etch. For all three systems, an total array aperture of 2σ to 2.5σ is enough to avoid truncation loss. The deeply etched AWGs show a more than 2 dB higher insertion loss than their shallowly etched counterparts. Part of this may be a calculation artefact, as explained above. When waveguide propagation losses are taken into account (figure 7.12(b)), a 2σ angular spread is an optimal choice with respect to insertion loss. Of course, this depends on the free spectral range of the device.

7.5.3 Crosstalk

As we will see in the next section, a lack of isolation between channels currently is the greatest bottleneck. As can be expected from the star coupler simulations, truncation crosstalk can form a large contribution. Secondly, we expect phase errors in the arrayed waveguides to generate a lot of crosstalk, just as with the Mach-Zehnder lattice filters in chapter 6. Finally, coupling between the arrayed waveguides can contribute to crosstalk. However, this latter source was not properly simulated here.

As we have seen, the minimum angular spread of the array will be determined more by the crosstalk due to truncation rather than by the insertion loss it generates. Therefore, the crosstalk for an ideally fabricated AWG was calculated as a function of angular spread. Again, 9-channel AWG router devices were simulated with constant parameters except for the angular spread and thus the number of waveguides. The devices were only calculated on their (real) channel wavelengths. As normally the pass-

(a) deep versus shallow etch, w_a=1.5 μm

(b) Influence of waveguide loss on a deep-shallow AWG

Figure 7.12: Simulated insertion loss for a 9-channel AWG router with 200 GHz channel spacing, as a function of angular spread of the array. w_p = 1.5 μm, R =58.7 μm, ΔL=40.3 μm, g_a = 200 nm, $g_{p,min}$ = 600 nm.

bands of neighbouring outputs are generally not located on a maximum of a sidelobe in the transmission spectrum, the values plotted therefore do not relate to the maximum sidelobe levels of the spectra, but rather to the actual crosstalk generated in the communication channels.

For the 200 GHz AWG of the previous paragraph, the crosstalk is now plotted against angular spread of the array in figure 7.13 for different values of the minimum port spacing. One artefact clearly appears: the finite overlap between the port waveguides generates a crosstalk floor for large angular spread (no truncation crosstalk). This is artificial, as we did not first propagate the isolated waveguide modes to the aperture. In reality, as long as the AWG is symmetric, coupling at the input side is inverted at the receiver side. With phase errors, the situation changes and coupling at the receiver side can have a real influence. The crosstalk floor is higher for shallowly etched ports, as expected. For deeply etched ports, the crosstalk has a minimum around a 3σ angle, with a 30 dB better isolation compared to a 2σ angle for large enough port spacing (low enough overlap in the simulation here). For the shallowly etched ports, the optimum angle is around 2σ. The *wavy* behaviour of the curves is real: as more waveguides are added, the filter order increases and thus more zeros are added. Therefore, more but denser spaced sidelobes appear in the spectrum. When scanning over the number of waveguides, the frequency channels move
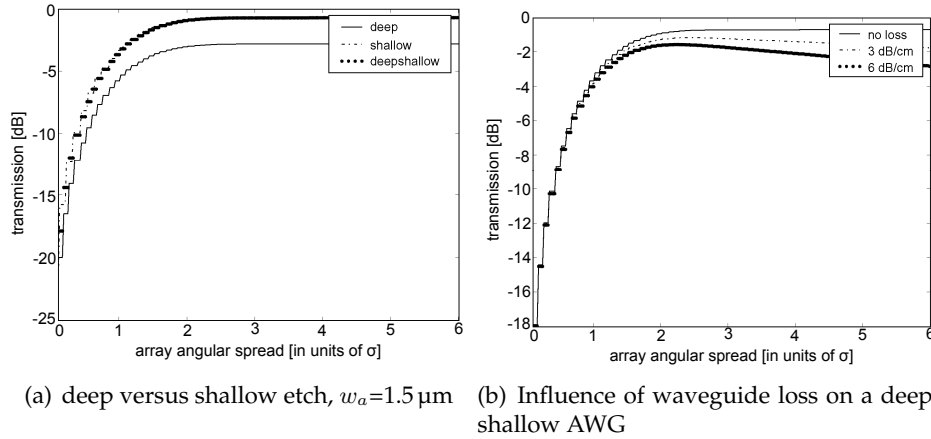
(a) deep etch

(b) deep-shallow etch

(c) shallow etch

Figure 7.13: Simulated crosstalk for a 9-channel AWG router with 200 GHz channel spacing, as a function of angular spread of the array, for varying minimum port spacing (varying focus). $w_p = w_a = 1.5\,\mu m$, ΔL=40.3 μm, $g_a = 200\,nm$.

over minima and maxima of sidelobes, generating the typical ripple in the plots.

Figure 7.5.3 checks whether propagation loss in the arrayed waveguides can have an influence on the crosstalk performance. Clearly, for the propagation loss levels achieved, this is not the case, at least for the relatively wide FSR (small delay length) in the example.

To simulate the influence of phase errors in the arrayed waveguides, random phase errors were added through the model described in appendix A. The crosstalk was calculated for 100 samples. Figure 7.15 plots the mean crosstalk along with its standard deviation indicated as vertical error bars, for 100 GHz and 200 GHz channel spacing. The waveguide width deviations correspond with what could be realistic values, as seen in chapter 6.
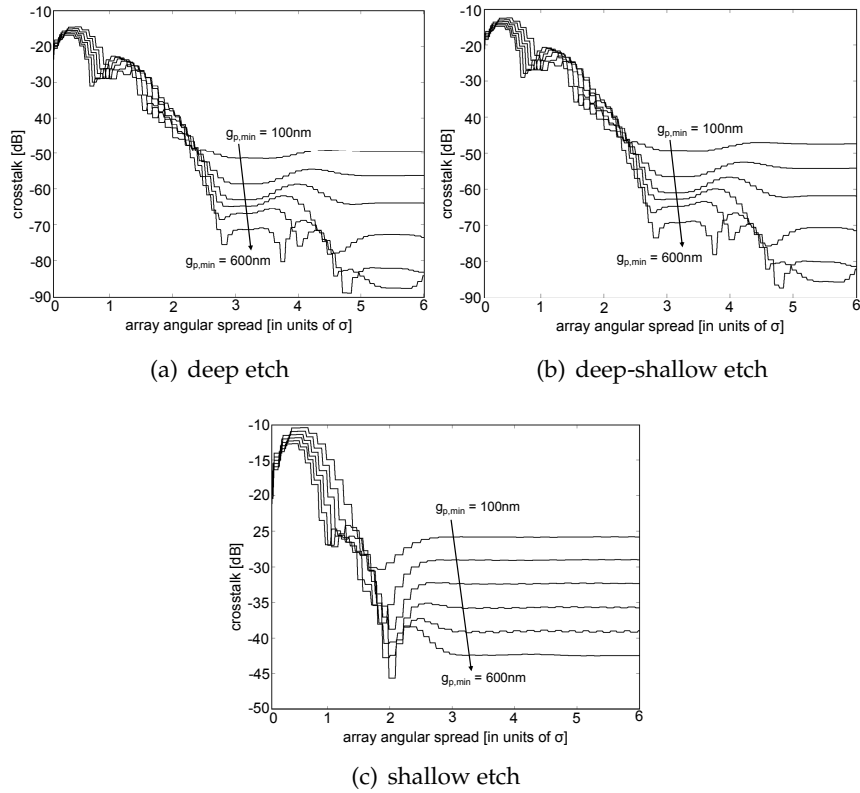
Figure 7.14: Simulated influence of waveguide loss on crosstalk for a 9-channel AWG router with 200 GHz channel spacing, as a function of angular spread of the array. Deep-shallow etch. $w_p = w_a = 1.5\,\mu m$, ΔL=40.3 μm, $g_a = 200\,nm$.

The crosstalk is higher for smaller channel spacing due to the longer delay length (smaller FSR). In the case of this 9-channel AWG, the difference is limited to a few dB. Clearly, the optimum angular spread of the array is defined by the phase errors in the case of high phase error. However, the insertion loss dictates an angular spread of around 2σ, if waveguide loss is incorporated. We can quickly see that even for only 2 nm width deviations on a 500 nm length scale, the effect is dramatic. With larger width variations of the order of 5 nm, the AWG device will not scale anymore to more than eight ports in a real-life application.

Clearly, for the phase error levels expected from the simulations and measurements in chapter 6, the crosstalk level is extremely high. Here we only simulate width variations, but from chapter 3, we can also expect small silicon layer thickness variations between different areas in an AWG, as the devices quickly grow a few hundred micrometre long in at least one direction.

7.5.4   Crosstalk reduction through waveguide broadening

Just like in chapter 6, a partial solution can be found in making the arrayed waveguides broader in order to reduce the dependency of the effective index on the waveguide dimensions. In the bends, the waveguides need to be single mode, and should therefore be kept narrow. This was implemented in the AWG mask designs and simulation code. Short linear tapers convert between the narrow and broader wires, as illustrated further on in figure 7.22(a). When we now look at the simulated crosstalk with such expanded waveguides for the same width variations as in the

(a) Δf_{ch}=100 GHz

(b) Δf_{ch}=200 GHz
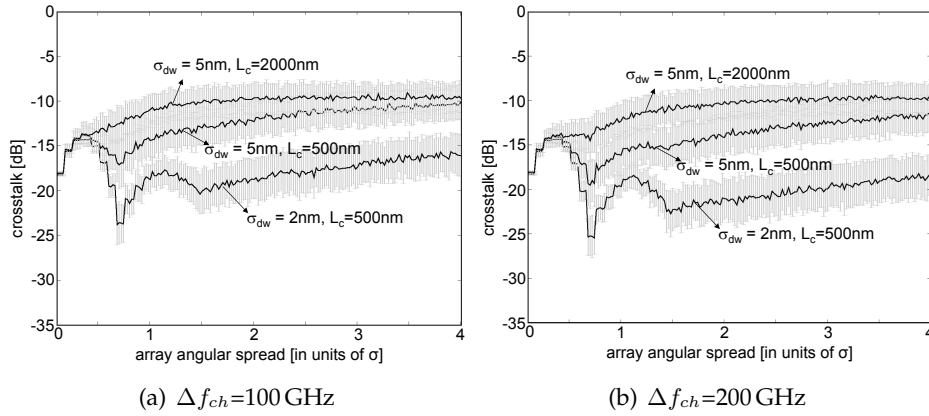
Figure 7.15: Simulated influence of phase errors on crosstalk for a 9-channel AWG router, as a function of angular spread of the array. Deep-shallow etch, for 200 GHz and 100 GHz channel spacings. $w_p = w_a = 1.5\,\mu m$, ΔL=40.3 μm or 80.6 μm, $g_a = 200\,nm$. Mean crosstalk is plotted, gray bars indicate standard deviation.

previous paragraph, given in figure 7.16, crosstalk is clearly lower. For a 100 GHz channel spacing the improvement in isolation is around 5 dB for a 2σ angle spread of the array. For larger channel spacing, the gain is smaller of course. A 5 dB improvement might not seem a great accomplishment, but in fact this means a serious reduction of the phase errors. Despite the waveguide broadening, waveguide width still needs to be controlled to the nm-level in order to obtain a crosstalk that makes the component scalable to the order of 16 or 24 channels. The scalability of an SOI AWG will be discussed in the context of an optical interconnect application in chapter 9.

7.5.5 Crosstalk as a function of number of channels and phase errors

If the processing technology is optimized so that waveguide width variations are minimized, phase errors can be seriously reduced. Figure 7.17 plots the mean expected crosstalk as a function of the correlation length L_c and the standard deviation σ_{dw}, for 100 GHz and 200 GHz channel spacings. In order to reduce the crosstalk caused by phase errors to lower than -35 dB, σ_{dw} should be reduced to at most 1 nm, and simultaneously the cor-

(a) Δf_{ch}=100 GHz (b) Δf_{ch}=200 GHz
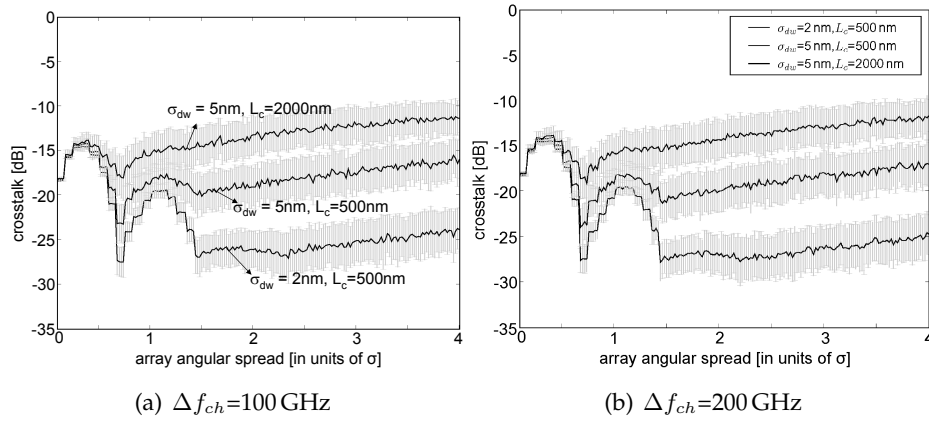
Figure 7.16: Simulated influence of phase errors on crosstalk for a 9-channel AWG router with expanded waveguides, as a function of angular spread of the array. Deep-shallow etch, for 200 GHz and 100 GHz channel spacings. $w_p = w_a = 1.5\,\mu m$, ΔL=40.3 µm or 80.6 µm, $g_a = 200$ nm. Waveguides are expanded to 800 nm width. Mean crosstalk is plotted, gray bars indicate standard deviation.

relation length should be 200 nm at most. The difference between 100 GHz and 200 GHz spacing is small, at least for eight channels.

To see the influence of the number of frequency channels on crosstalk, the mean crosstalk was calculated for different channel counts, for an AWG with an angular spread of 2σ. The result is plotted in figure 7.18 for different levels of phase errors. For 4 to 32 channels, the variation of crosstalk is only a few dB. The crosstalk level is rather defined by the phase error level. However, a higher channel count necessitates a lower crosstalk level, so that the achievable number of channels with current technology is limited.

(a) $\Delta f_{ch} = 100\,\text{GHz}$

(b) $\Delta f_{ch} = 200\,\text{GHz}$

Figure 7.17: Simulated influence of phase errors on crosstalk of an AWG router as a function of correlation length L_c and σ_{dw}. Deep-shallow etch. Angular spread of the array: 2σ. $w_p = w_a = 1.5\,\mu\text{m}$, $g_a = 200\,\text{nm}$. Mean crosstalk is plotted.

Figure 7.18: Simulated influence of phase errors on crosstalk of an AWG router for different channel counts. Deep-shallow etch, for $100\,\text{GHz}$ channel spacing. Angular spread of the array: 2σ. $w_p = w_a = 1.5\,\mu\text{m}$, $g_a = 200\,\text{nm}$. Mean crosstalk is plotted, gray bars indicate standard deviation.

7.6   Experimental results

A large number of AWG devices were fabricated, with varying parameters, but all with 100 to 400 GHz channel spacings. We will cover the most notable ones in the sections below, in chronological order. All devices reported on were designed as routing devices, which intrinsically limits their insertion loss non-uniformity. Some deeply etched star couplers were characterized as well, which will be covered first.

7.6.1   Design procedure

All designed AWG devices were designed as router devices with a cyclic behaviour. Therefore, the delay length is set from the required FSR, and the insertion loss non-uniformity is intrinsically of the order of 2 to 3 dB. A slightly different design procedure was used than described in section 7.2. The number of waveguides N is given. The resulting procedure is

- calculate ΔL and filter order from the desired FSR and central wavelength

- the angular spread of the array is calculated from the width of the port aperture

- the focus distance R is calculated from the array pitch, the number of waveguides and the angular spread.

- the port angles are computed from the focus distance and the dispersion.

A lot of software was written to put the different types of AWG devices studied on masks. The parameter calculation code was integrated with the mask design code. Both were written by Dr. W. Bogaerts, who designed (almost) all AWG mask structures.

7.6.2   Star couplers

Deeply etched star couplers were fabricated, with the output waveguides covering a large angle. Only 16 of the output waveguides were actually connected to fibre couplers and the transmission to these was characterised. Three star couplers were designed, with 1.0 μm, 1.5 μm and 2.0 μm wide apertures, respectively.

Vertical fibre couplers were used, making the measurements reliable in a 1530 nm to 1555 nm range for the outer waveguides receiving only

a small fraction of the power. The measured power transmission is plotted in figure 7.19 for a port width w_p of 1.0, 1.5 and 2.0 μm, as a function of the angle the output waveguide is mounted on. The measurement for w_p=1.5 μm coincides very well with the simulated transmission for a 100 nm port width. This simulation was done through the method described in section 7.4.2, using a horizontal slice through the 3-D mode profile. The same holds for the 2000 nm port width, apart from a 2 dB difference in power. For the 1000 μm port width, the simulation underestimates the tails and thus the truncation loss and crosstalk. The difference between simulation and measurement may be explained by the sidewall slope of the waveguides.

Table 7.2 shows the measured diffraction angle of the star couplers, and the power in the far field tails, for several cutoff power levels. As only 16 outputs were measured, a cubic spline interpolation was done first. The wider waveguides show a clearly lower diffraction angle, but the power in the 2σ to 3.5σ tails is comparable. In an AWG the tails are lost and generate loss, but also at most half of the truncated power is coupled into the neighbouring output ports as crosstalk. In table 7.2, this maximum nearest-neighbour crosstalk is given too. A limit for AWG devices in SOI clearly appears here. If the array acceptance angle is chosen as the 2σ diffraction angle, then the impact on overall loss would be limited, but the crosstalk can already be as high as -16.3 dB, without any other crosstalk contributions taken into account.

In order to accept most of the tail power and push the maximum crosstalk to the -30 dB to -40 dB range, the opening angle of the array should be as high as 80°to 90°for reasonably wide waveguides. The total array angle should be larger than the 3σ angle, ideally almost 4σ, in order to keep crosstalk within bounds. In reality, an optimum angular spread of the array will be defined by this decreasing crosstalk with increasing angle, and the increasing crosstalk due to phase errors in the longer waveguides.

These figures provide a good reason to choose the port aperture width on the wider side, at least 1.5 μm. Alternatively, the ports can also be implemented with a shallow etch, exhibiting a smaller diffraction angle due to the enlarged spot size. Star couplers with a shallow port or arrayed waveguide etch were not characterised, however. It should be noted that the crosstalk figures calculated are a worst-case scenario, and will in reality be better.

Choosing a large port aperture width will increase the star coupler length. On the other hand the port spacing can stay the same or smaller,
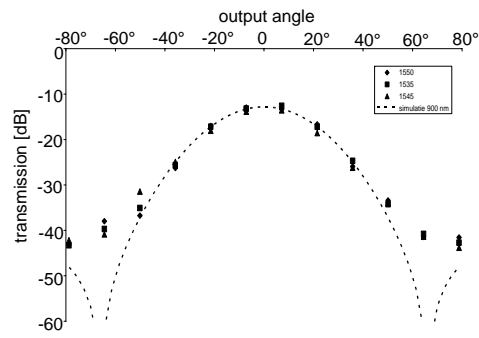
(a) aperture width 1.0 μm

(b) aperture width 1.5 μm

(c) aperture width 2.0 μm

Figure 7.19: Measured transmission of star couplers with deep etch. From central input port to all output ports connected to fibre couplers, for three wavelengths. Dotted line: simulation with 100 nm narrower aperture than designed.

property/aperture width	1.0 μm	1.5 μm	2.0 μm
diffraction angle @ 2σ	57°	39°	30°
truncation loss @ 2σ	0.047 dB	0.047 dB	0.037 dB
max. eq. crosstalk @ 2σ	-16.3 dB	-16.3 dB	-17.3 dB
diffraction angle @ 3σ	90°	57°	41°
truncation loss @ 3σ	0.0054 dB	0.0035 dB	0.006 dB
max. eq. crosstalk @ 3σ	-25.6 dB	-27.6 dB	-25.3 dB
diffraction angle @ 3.5σ	128°	81°	63°
truncation loss @ 3.5σ	0.001 dB	0.002 dB	0.001 dB
max. eq. crosstalk @ 3.5σ	-32.8 dB	-30.9 dB	-31.7 dB

Table 7.2: Measured transmission through star couplers, from centre input. Full (two-sided) diffraction angle, power in the truncated tail, and maximum nearest-neighbour truncation crosstalk, for 2σ, 3σ and 3.5σ angular aperture.

as broader waveguides exhibit less coupling, and the passband width will increase with port width.

7.6.3   First generation

The first AWGs made were deeply etched overall, had no broadened waveguides, and used the type of layout common to lower index contrast devices.

The transmission of an eight channel AWG with a designed channel spacing of 400 GHz is shown in figure 7.20. The layout of the AWG was classical, with a large bend radius of 75 μm. The starcouplers were etched deeply. The measured on-chip loss is as high as 9.5 dB, and the crosstalk of this example is -7 dB. Other devices from the same mask show even higher crosstalk, up to transmission spectra where the main lobe cannot be discriminated from the sidelobes. The high loss has three causes: scattering at the gaps between the deeply etched array apertures, loss to other diffraction orders, and an underestimated diffraction angle. The insertion loss for this AWG decreases slightly to 8 dB with a lower illumination dose during lithography, giving more narrow gaps and broader waveguides.

The high crosstalk seems to bear no relation to the aperture spacing, as other AWG devices with larger output port spacings experience a larger crosstalk. There is no influence of the lithography illumination dose on the crosstalk level. However, the underestimated diffraction angle can give rise to a nearest-neighbour crosstalk as high as -10 to -11 dB. Phase errors

in the arrayed waveguides superimposed on this, can quickly drive the crosstalk as high as the measured levels.

7.6.4   Second Generation

For the second generation, AWG devices were designed with shallow etched array and port apertures, in order to reduce the insertion loss. Unfortunately, the diffraction angle was underestimated with a factor of $\sqrt{2}$ [8]. Four four-channel devices with shallow etch were designed, two with 200 GHz and two with 250 GHz channel spacing. Of each type, one AWG had 10 waveguides and another one 14. Otherwise, the devices were identical, with aperture widths of 1.5 µm and a gap of 0.2 µm.

The devices with shallow etch were characterised through an eight fibre array connector, as described in chapter 9. The exact insertion losses are difficult to extract, but in chapter 9, the on-chip insertion loss was derived to be maximally 3.5 dB for the inner channels. This can also be seen from the transmission spectra for the outer channels, given in figure 7.21. This is a clear reduction compared to the 8 dB loss shown by deeply etched AWG devices. Such deeply etched devices were again incorporated on the mask with the second generation, but all showed very high crosstalk, with some outliers in the positive sense (-5 to -7 dB).

The crosstalk was generally high, but all four devices showed a reduced crosstalk at a lower, but not always the same, lithography dose (wider waveguides). Otherwise, there was no clear dependence on lithography dose. Figure 7.21 shows the best performing device, with a sidelobe level reaching -12 dB. Interestingly, other nominally identical samples showed a maximum sidelobe level of the same order, but different sidelobe forms.

7.6.5   Third Generation

In a third generation of AWG devices, two problems were tackled. First, the total array angle was corrected. Still, only the 2σ angle was used, which in a worst-case limits the nearest neighbour crosstalk to around -17 dB. Based on measured star coupler data for deep star couplers, the difference in insertion loss compared to the second generation AWGs can be around 1.4 dB. The largest contribution to the sidelobes is thought to be the phase errors between the arrayed waveguides, so the second change was to expand the waveguides to 800 nm in order to reduce the influence

---

[8]Due to a misinterpretation of the simulation software results

(a) Microscope view

N_{arms}	18
Bend radius	75 µm
$w_{i/o}$	1.0 µm
g_i	0.2 µm
g_o	0.1 µm
w_a	1.0 µm
g_a	0.18 µm
total array aperture angle	63°

(b) Parameters

(c) Transmission

Figure 7.20: First generation: 5×8 AWG with $\Delta f_{ch} = 400\,GHz$ designed and deeply etched star couplers. Microscope picture and transmission from central input the all eight outputs.

(a) Mask layout. The align ports are connected two by two
(1,8), (2,7) and (3,6)

(b) Input 1

(c) Input 4

Figure 7.21: Second generation: AWG router with double etch, 4 channels,
$\Delta f_{ch} = 250\,GHz$, $N_{arms} = 10$, $w_{i/o/a} = 1.5\,\mu m$, $g_{i/o/a} = 0.2\,\mu m$. Measured fibre to fibre transmission spectra from two inputs to all four outputs.

name	exp. wg	Δf_{ch} [GHz]	N_{ch}	N_{arms}	$w_{i/o/a}$ [μm]	g_a [μm]	$min(g_{i/o})$ [μm]	R [μm]
Karl	yes	100	8	16	1.5	0.2	0.455	49.2
Gummo	yes	400	8	19	1.5	0.2	0.760	58.6
Harpo	no	400	8	19	1.5	0.2	0.750	58.5
Chico	yes	400	6	15	1.2	0.2	0.925	33.8
Groucho	no	400	8	19	1.5	0.2	0.760	58.6
Zeppo	yes	200	16	36	1.5	0.2	0.660	111

Table 7.3: Parameters of the designed AWG devices on the PICMOS02 mask (third AWG generation). All AWGs were designed as routers. Every device was assigned a name.

name	max sidelobe at best dose	wire width 450 nm designed	wire width 800 nm designed
Karl	-9 dB	532 nm	840 nm
Gummo	-12 dB	575 nm	890 nm
Chico	-9 dB	570 nm	886 nm
Zeppo	< -16 dB	560 nm	875 nm

Table 7.4: Third generation: Measured crosstalk performance of the AWG devices with expanded waveguides

of small width and height variations. In the bends, the waveguides were kept narrow to prevent mode coupling. Note that enhancing the array angle in order to reduce truncation crosstalk, will enlarge the phase errors. Therefore, the net effect may be a saturation or an increase of the crosstalk level, as seen from the simulations in the previous section.

The parameters of the designed routers are listed in table 7.3. As the routers are designed with a constant frequency spacing, the input and output waveguide pitch is slightly varying. Therefore, the minimum pitch is given. All routers designed with confocal star couplers.

As a single mask contains many components, and over multiple reticles, a large number of components and the accompanying measurement data are accumulated, we started naming the devices with easily callable names. This avoids confusion due to numbering components, needing to count structures on a sample, or difficulty remembering component characteristics after some time.

Harpo and Groucho, the two almost identical devices without expanded waveguides, performed very similarly, with for most lithography doses a spectrum with unrecognisable transmission bands. On the other hand, the other devices with expanded waveguides all and always showed a clear AWG spectrum, albeit with sometimes high sidelobes. The crosstalk performance of the AWGs with expanded waveguides is given in table 7.4. The AWG with the 100 GHz channel spacing and smallest FSR shows a higher nearest-neighbour sidelobe at around -9 dB. However, the other sidelobes are lower, at around -12 dB to -14 dB. The other devices don't show this nearest-neighbour behaviour or to a lesser extent (in Zeppo). The simulated maximum sidelobe levels of Karl, Gummo and Chico without phase errors are -32 to -34 dB.

Zeppo, the 16 channel AWG with 200 GHz spacing, shows a consistently better behaviour, even though the number of waveguides is significantly higher at 36. A SEM picture and transmission spectrum of the device are shown in figures 7.22(a) and 7.22(b). The crosstalk is better than -16 dB, reaching -20 dB for some channels. As the delay length is the same as for the eight channel devices with 400 GHz channel spacing, the difference in crosstalk level cannot be attributed to phase errors. The insertion loss of this device is around 2.2 dB, which is an improvement compared to the second generation AWG devices which can be attributed to the wider total angle of the array. The simulated transmission is shown in figure 7.22(c), showing good accordance. The measured passbands are somewhat narrower, possibly due to narrower than designed port apertures. The measured insertion loss is higher, and the simulation underestimates the insertion loss non-uniformity to a limited extent. The simulation shows that overlap between the passbands already limits the isolation to -25 dB. With 36 waveguides and 16 output channels, this AWG still has a footprint of only 0.1 mm^2.

As a conclusion for this third generation of devices, the crosstalk level was reduced. While in the second generation an isolation of around 10 dB was rather good luck, this is now made the expected performance through the broadening of waveguides to 800 nm. One notably good design showed an isolation better than 16 dB. However, further experimental data on the influence of several parameters is needed, to see what is limiting the crosstalk.

7.6.6 Fourth generation

A fourth mask (PICSOI08) contained sets of AWG devices to test for several influences:

(a) SEM pictures

(b) Measured transmission (central input to all outputs)

(c) Simulated transmission (central input to all outputs)

Figure 7.22: Third generation: a 16 channel AWG with 36 arrayed waveguides, $\Delta f_{ch} = 200\,\text{GHz}$, and a 0.1 mm^2 footprint

Name	Δf_{ch} [GHz]	etch ports	etch arms	N_a	rotation	base length [µm]	expansion
Porthos	200	shallow	shallow	18	90°	+0	gen.3
Athos	200	deep	shallow	18	90°	+0	gen.3
Aramis	200	deep	deep	18	90°	+0	gen.3
Mickey	200	shallow	shallow	18	90°	+0	gen.4
Donald	200	deep	shallow	18	90°	+0	gen.4
Goofy	200	deep	deep	18	90°	+0	gen.4
Shiva	400	shallow	shallow	10 to 34	90°	+0	gen.4
	400	shallow	shallow	26	90°	-40	gen.4
						+100	
						+250	
	400	shallow	shallow	34	70°	+0	gen.4
					45°		
					0°		
Vishnu	400	deep	shallow	10 to 34	90°	+0	gen.4
	400	deep	shallow	34	70°	+0	gen.4
					45°		
					0°		

Table 7.5: AWG designs of the fourth generations. All devices have 8 channels and are designed as cyclic routers. The port and arrayed waveguide aperture widths are 1.5 µfor all devices, and the arrayed waveguide aperture gap was designed 200 nm. Each AWG was designed a name for easy reference.

- Deep versus shallow etch for port and array apertures.

- The expanded waveguides from generation 3, versus a corrected algorithm to have all the delay length in expanded waveguides.

- The angular spread of the array.

- The base length of the array waveguides.

- Rotation of the AWG on the reticle/wafer.

Table 7.5 lists the AWGs designed. We will discuss each of the influences measured. In this way, we can hope to see whether truncation, phase errors, or other errors (due to waveguide coupling for instance) are limiting the crosstalk.

Figure 7.23: Measured maximum sidelobe of AWGs with different number of waveguides (spreading angle), on two samples (from columns 0 and -4 on the wafer) with different lithography dose.

Angular spread of the array

To test the influence of the angular spread, the number of waveguides of an AWG was varied, while keeping the length of the inner waveguides constant. (Normally, an AWG with a smaller number of waveguides also has shorter waveguides). All other parameters were kept constant. The resulting crosstalk levels are indicated in figure 7.23. The sidelobe level decreases as the number of waveguides grows due to smaller truncation, up to some optimal level, and then decreases again for a higher number of waveguides, as the aggregate phase error grows. For 3 of the 4 devices (two sets on two samples with different lithography doses), 30 waveguides was an optimum, while one device showed the highest extinction for 26 waveguides. This is for an angular spread of 3σ and 2.6σ of the calculated far field intensity respectively, with the 1σ angle as defined in section 7.5.1 This corresponds with conclusions made from the measurements of deeply etched star couplers in section 7.6.2. The best crosstalk levels reached are -14.4 dB and -14.66 dB, for Shiva (shallow port and array apertures) and Vishnu (deep port and shallow array apertures), respectively. However, for the latter AWG, the ones with 30 and 34 arms clearly show a nearest-neighbour behaviour, while the one with 26 arms does not. The 30 arm AWG showed the lowest overall sidelobes. Therefore, using a higher number of arms may still be beneficial.

If we look at the crosstalk at the channel wavelengths only, then the situation is different. This is shown by the simulation result in figure 7.24. For small enough phase error, the crosstalk stays constant from a certain point number of waveguides on, where crosstalk due to truncation becomes negligible. Analysis of the measurements confirms this behaviour: from 18 or 26 waveguides on, crosstalk at the channel wavelengths stays approximately constant, to within the simulated deviation of the crosstalk. However, the simulated crosstalk only shows this behaviour for lower phase errors that cannot explain the crosstalk levels measured. This discrepancy may have two causes. First, it is possible that the approximation used underestimates the mode tails and high spatial frequencies of the real fields and thus the crosstalk. Secondly, it may indicate that other mechanisms causing crosstalk have a significant contribution and are not depending as much on the angular aperture or waveguide length, or that these other mechanisms are superimposed on the influence of phase errors in the waveguides due to width deviations in a non-straightforward way. Possible sources of crosstalk in the current structures include interference with scattered light between both star couplers [87], coupling between the arrayed waveguides (both in the tapers and the waveguides as such), distributed scattering in the waveguides as evident form chapter 5, and coupling between the port apertures combined with imperfect imaging.

It should be noted that the number of devices measured is small however, so the conclusions are only indications and not based on statistically relevant data.

Two example transmission measurements of Shiva (shallow) and Vishnu (deep-shallow) with 30 arrayed waveguides are given in figure 7.25, taken from the same sample. Note that with this larger number of arrayed waveguides, the improvement in crosstalk is traded for a decrease in passband width.

Base length of the array

The AWG Shiva has a delay length of 24.5 μm. For 26 arrayed waveguides, corresponding to a 2.6σ angular aperture, the base length of the array was decreased or increased to see the influence of phase errors. When 100 or 250 μm is added to each array waveguide of shallowly etched AWG with 26 waveguides, performance deteriorates to sidelobe levels higher than -2.2 dB, on devices on two different samples. The variance of the phase error is high. For the waveguides of medium length, carrying the bulk of the optical power, the length increase is important (of the order of 25% to >100%), leading to a $\sigma_{d\phi}$ of up to 50% larger than for the standard base

(a) Shiva (shallow etch)

(b) Vishnu (deep-shallow etch)

Figure 7.24: Simulated maximum crosstalk of AWGs as a function of spreading angle, for different phase error levels.

(a) Shiva (shallow etch)

(b) Vishnu (deep-shallow etch)

Figure 7.25: Measured mean crosstalk from the centre input to all 8 outputs for 8-port AWGs with 30 arrayed waveguides, for shallow and deep-shallow etch. The eighth output port was included in an interferometer and therefore has an excess loss of at least 3 dB. Gray bars indicate standard deviation.

Figure 7.26: Mean crosstalk at the channel wavelengths for an eight-port AWG with varying minimum waveguide length. Shallow etch (Shiva). $\sigma_{dw} =$ 5nm,L_c=500nm. Grey bars indicate standard deviation.

length. This increase of crosstalk is also predicted by the simulation result shown in figure 7.26, however at lower phase error and crosstalk levels.

For the standard base length and a 40 μm shorter length, crosstalk is -8.9 down to -14.4 dB. However, due to the variance of the phase error, the 40 μm shorter array does not consistently exhibit a better crosstalk performance. This relation is confirmed by figure 7.26. This indicates the main source of crosstalk still is the stochastic phase errors in the arrayed waveguides.

Rotation on the reticle/wafer

The rotated Shiva and Vishnu AWGs were measured on two samples with different lithography doses. An overview is given in figure 7.27. Three of the four devices show the same behaviour. The AWG rotated under 90 °, so oriented along the east-west axis on both reticle and wafer, showed the lowest sidelobes, at -11.9 to -14.7 dB. Rotated under 70 °, the crosstalk was higher, and under 45 °even higher. Placed along the north-south axis, the crosstalk performance was better again, but not as high as under

Figure 7.27: Maximum sidelobe of identical AWGs rotated under varying angle, on two samples (0 and -4) with different lithography dose.

90 °rotation. Vishnu does not show this behaviour on sample 0, however. It may be just an outlier. This may point either to an angle-dependent variation on the reticle, or to a directional behaviour of the fabrication process. For instance, there may be a slight dependence of the etching process on the crystal axes. The source may be identified by rotating the mask with respect to the wafer (at the cost of a wafer and processing run).

Deep versus shallow etch

The influence of the etch depth and the index contrast around the star couplers is not straightforward. With a lower index contrast, the minimum port spacing $d_{p,min}$ should be larger in order to avoid crosstalk due to coupling if imaging is not perfect. However, if this is done properly, there is no clear difference in crosstalk between a lower or higher index contrast at the port side, as discussed already, but the downside of a lower index contrast is a larger device size. With the design scheme used here, with N_{arms} given and $d_{p,min}$ defined in the process, the port spacing is automatically corrected for index contrast, as the angular spread of the array, and with given N_{arms} the angular pitch of the array, are calculated from the diffraction angle of the port apertures. For the AWG devices with shallow port etch (Porthos, Mickey), this resulted in a minimum port pitch of 680 nm, while the deeply etched ports had a minimum pitch of 245 nm.

(a) Shallow-shallow

(b) Deep-shallow
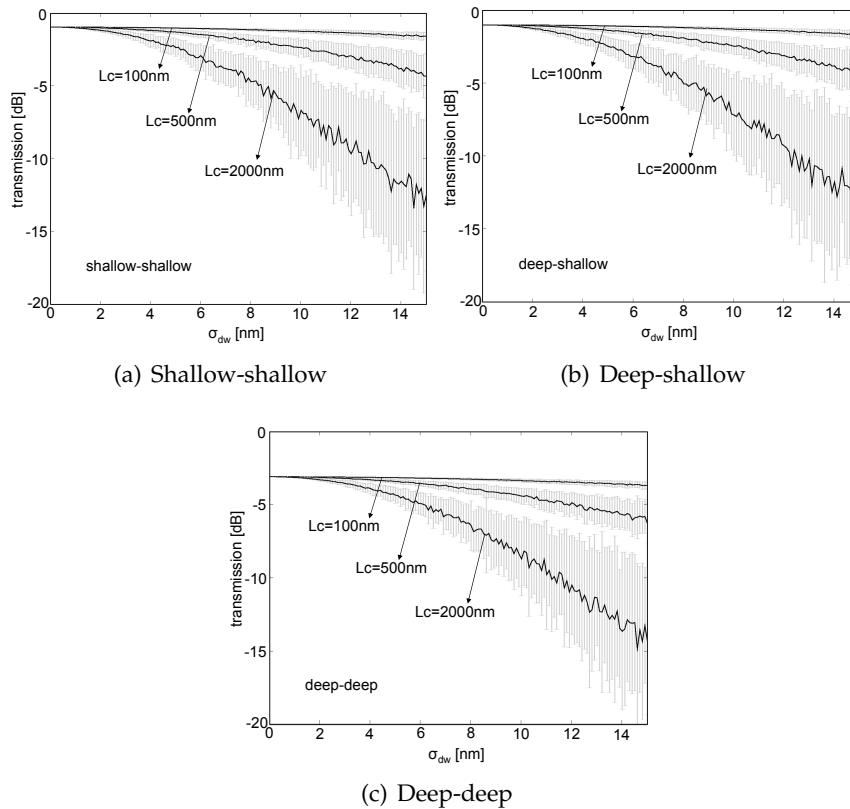
(c) Deep-deep

Figure 7.28: Simulated insertion loss of Porthos (shallow-shallow), Athos (deep-shallow) and Aramis (deep-deep), for a varying phase error.

Figures 7.28 and 7.29 plot the simulated insertion loss and crosstalk for the three AWG devices as a function of phase error. An insertion loss difference of a few dB is expected between shallow and deep etch for the array waveguides. For the crosstalk, the fully deeply etched AWG is expected to have the lowest crosstalk, although the difference is small.

The measurement results are given in figure 7.30. A difference between the gratings of the third and fourth generation can be seen. In the fourth generation, all delay length was implemented with expanded waveguides. Nevertheless, the mean insertion loss and crosstalk seem to be a couple of dB higher. This may be related to the extra pair of narrow to wide tapers in each arm. A difference between an oxide and air top cladding can be seen in the crosstalk, with the oxide clad devices performing better overall. Due to the large variations in loss and crosstalk due

(a) Shallow-shallow

(b) Deep-shallow
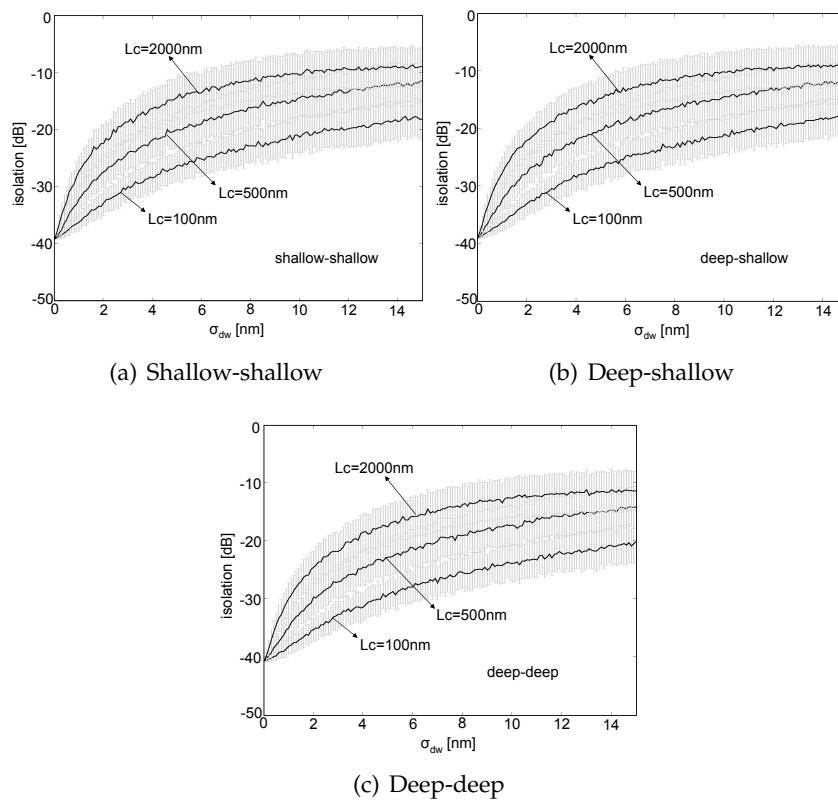
(c) Deep-deep

Figure 7.29: Simulated mean crosstalk of Porthos (shallow-shallow), Athos (deep-shallow) and Aramis (deep-deep), for a varying phase error. Gray bars indicate standard deviation.

(a) Insertion loss, gen. 3 grating

(b) Insertion loss, gen. 4 grating

(c) Crosstalk, gen. 3 grating

(d) Crosstalk, gen. 4 grating
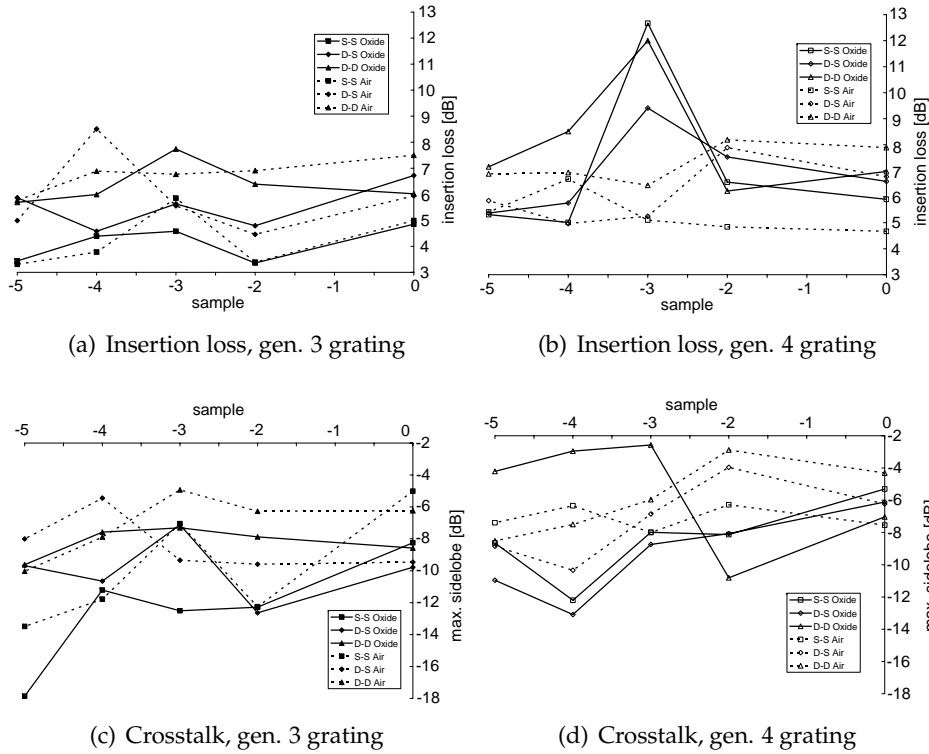
Figure 7.30: Insertion loss and maximum sidelobe of AWGs with shallow etch for both ports and array (S-S), deep etch for ports and shallow for array (D-S), or deep etch for both (D-D). On several samples with dose varying from low to high, with oxide and air claddings.

to high phase errors, and the limited number of measurement points, it is difficult to draw absolute conclusions from the data. Nevertheless, a trend can be seen. Overall, the fully deeply etched devices perform worst both in insertion loss and crosstalk. Crosstalk mainly comes from the nearest neighbours, while crosstalk from further off neighbours can be limited to the order of -15 dB. The fully shallowly etched devices perform best with respect to insertion loss. As expected, the difference in crosstalk between deeply or shallowly etched port waveguides is not clear, for shallowly etched array apertures.

A transmission measurement for the AWG with shallow etch and 200 GHz channel spacing (Mickey) is given in figure 7.31.

Figure 7.31: Measured transmission from the centre input to all outputs for an 8-port AWG with 18 arrayed waveguides and shallow etch (Mickey). The eight output port was included in an interferometer and therefore has an excess loss of at least 3 dB.

Conclusion

With the fourth generation of AWG devices, some parameters were scanned in order to get a better experimental view of the various influences on loss and crosstalk. Choosing an aperture angle beyond 2σ proved beneficial for the maximum sidelobe level, but did not influence the crosstalk at the actual channel wavelengths much. By making all waveguides shorter or longer, it was shown that phase errors are the major cause for the crosstalk. Nevertheless, the measurements, compared with the simulations, indicate that other factors have a significant contribution as well. The choice between deeply or shallowly etched port waveguides is yet undecided, as experimentally both show little difference with respect to loss and crosstalk.

7.7 Conclusion

Silicon-on-insulator AWG routing devices were demonstrated with 4 to 16 wavelength channels and 100 GHz to 400 GHz channel spacing. For simulation of the devices, computer code was written, integrating various approximation methods, material models and interfacing with mode solvers. By using a double-etch process to lower the index contrast around the star coupler, the insertion loss was reduced. Currently, an insertion loss of 2 to 3 dB is obtained. The phase errors in the arrayed waveguides, caused mainly by stochastic width variations, were reduced by broadening the waveguides. Nevertheless, crosstalk is still high, with isolation levels of 15 dB to 17 dB for the best AWG devices shown. Analysis of the measurement results indicate that while phase errors form the largest contribution to crosstalk, other sources have a significant contribution as well. The crosstalk might be optimized by changing the total aperture angle of the array. Simulation and experiments show that both with respect to loss and crosstalk at the actual channel wavelengths, a 2σ angle is optimal. The maximum height of the sidelobes in the transmission spectrum might be decreased by choosing a larger angle of around 2.7σ. The devices have a typical footprint of $0.1\,\text{mm}^2$, enabling many devices and extra functionality to be integrated on the same chip. Additionally, in chapter 9 it will be shown that already at the current level of performance, AWG components may be used for datacommunication applications.

Chapter 8

Reproducibility and accuracy

As discussed in chapter 2, small dimensional variations have a large impact on the effective index of a photonic wire. In chapter 3 a number of structural variations were described, such as sidewall roughness, layer thickness variation, mask digitisation and non-uniformity over the wafer. The channel isolation of AWG and lattice filters evidenced a significant level of dimensional variations. In this chapter, the reproducibility and accuracy of structures will be discussed further. Both non-uniformity over the wafer and non-uniformity within short distance (single die) are studied. Then we have a look at ways to improve on the reproducibility and accuracy or how to compensate for it.

From the characterisation of ring-resonators channel drop and all-pass filters, Mach-Zehnder lattice channel drop filters and arrayed waveguide grating devices in the previous chapters, we know already that significant dimensional variations are present within a component. However, on a longer length scale, deterministic and stochastic processes will give rise to a lack of reproducibility between components within the same die and between dies. Some different effects may exist between components in a single die or between dies:

- **Difference of the mean feature size due to a deterministic variation**. For instance, the resist thickness will vary over the wafer, mainly with a different thickness at the outer side of the wafer. This can influence both the lithography and etching process. Also, the etching process may not be perfectly uniform over the wafer, and the etch rate changes as a function of the structure density (*loading*).

- **Difference of the mean feature size due to stochastic variations within a wire**. Stochastic variation of the wire width within a single wire lead to a different mean width between wires on any length

scale. The stochastic variations may depend on the position on the wafer in some way, leading to a deterministic component in this variation of mean dimension. The different mean wire width between wavelength selective components leads to spectral differences between those components.

- **Difference of the mean feature size due to stochastic variations between wires**. In theory, some stochastic variation over the wafer could exist that gives feature size variations independent of the stochastic variations within a feature.

- **Different stochastic variations between components**. Not only lead the random variations to a different mean dimension and thus to a spectral shift, but each sample will have a different crosstalk signature. The larger the variance of the stochastic process, the larger the variance of crosstalk.

8.1 Characterisation

8.1.1 Within-die non-uniformity

To test the uniformity within a limited distance over the wafer, identical components were put on several positions on a die. The mask layout is given in figure 8.1. The structure consisted of a pair of identical all-pass ring resonators, at a distance of $dy = 25\,\mu m$ of each other, and a pair of identical Mach-Zehnders at a distance of $(dx = 99\,\mu m, dy = 12\,\mu m)$. This structure was repeated on the same die at a distance of $dy = 4650\,\mu m$ and $dy = 7375\,\mu m$. In this way, testing of short to longer scale uniformity is possible. Additionally, the structures were characterised on samples with different waveguide widths and gap widths obtained using different lithography illumination doses. The ring resonator has a radius of 4 µm and a coupling length of 4 µm, or a total unit delay length of 33 µm. The MZI has a delay length of 50 µm. An example of resulting measured spectra for the smallest waveguide width is given in figure 8.2.

From this data, it is clear large differences between nominally identical structures exist. For wider waveguides, we can expect the deviations to be smaller. It is already apparent the structures close to each other experience a smaller relative shift than structures spaced further apart.

In order to look at the differences on a short length scale, figure 8.3 plots the shifts between closely spaced filters for each of the three groups of identical structures within a die. Two features become apparent. First,

Figure 8.1: Mask layout for characterization of uniformity within a die. Two pairs of identical Mach-Zehnder interferometers and all-pass ring resonators close to each other are repeated on the die over a distance of several mm.

(a) All-pass rings

(b) MZIs
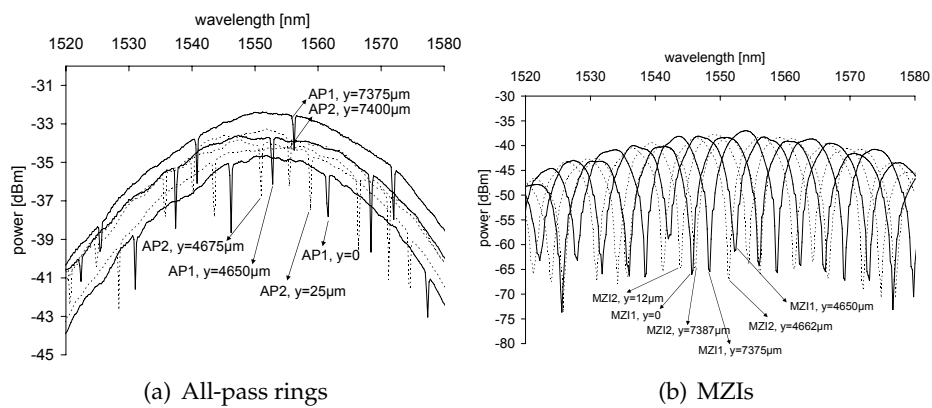
Figure 8.2: Within-die uniformity: Measured spectra of identical ring resonators and Mach-Zehnder interferometers placed near to each other (AP1,AP2 and MZI1, MZI2) on three locations of a die

the shift between the closely spaced filter pair that was designed identical, almost everywhere is negative, or below +500 pm. This indicates a deterministic component, which can be related to the two structures being different on the reticle, both for the ring and the Mach-Zehnder structure. The origin of this difference may already lay in a difference in the mask design, or may be due to fabrication variation of the reticle. Another explanation may be a dependence of the processing on the North-South direction along the wafer. Secondly, for both the ring and MZI filter, the shifts become larger for more narrow waveguides, which is to be expected, as small width variations have a larger influence for narrower waveguides.

Quantitatively, the short-distance spectral non-uniformity does not surpass 3.5 nm, and is generally of the order of 0.5 nm to 2 nm. These values are related to a change of the mean waveguide width in figure 8.4, using the values calculated in chapters 2 and 4, supposing no other variations. The DC component was subtracted from the data in order to compensate for the deterministic component, supposedly constant with respect to wire width.

In conclusion, concerning the short-distance non-uniformity, there is a deterministic component between two components spaced 12 to 25 μm apart, which can be related to a mean width difference of 1 to 2 nm, and stochastic variations superimposed on that have a standard deviation of the mean width below 1 nm. For narrow waveguides, the spectral shift is of the same order, while for wider waveguides, the spectral shift is slightly better. Note that these width changes are far below 1%. Improving this with CMOS processes adapted for fabrication of photonic structures of various nature will prove rather impossible.

Using the model in appendix A, we obtain an idea of the order of magnitude of the stochastic width variations that can result in these mean width variations. (See appendix, figure A.2). For a wire length of 10 to 100 μm, the typical delay length in the filters studied, a correlation length in the 500 to 1000 nm would be needed to explain a $\sigma_{\mu(dw)}$ of 0.5 to 1 nm, and a $\sigma_{d(sidewall)}$ of at least a few nm. So, in terms of the stochastic width variation, the $\sigma_{dw} = \sqrt{2\sigma_{d(sidewall)}^2}$ would be of the order of 4 to 6 nm, with a L_c of 500 to 1000 nm. From the SEM pictures in chapter 3, we can see this is very well possible.

In figure 8.2, we already see that the non-uniformity over a longer length scale is larger. By taking the mean of the stopband centres of each identical filter pair, and comparing the pairs separated apart on the order of a few mm, we arrive at figure 8.5. Spectral shifts of up to ±10 nm are obtained. Similar to the short-distance shifts, these can be related to

(a) All-pass rings

(b) MZIs

Figure 8.3: Within-die uniformity: differences between stopbands of filters (AP1,AP2 and MZI1,MZI2) at short distance of each other, as a function of isolated wire width. The minimum, median and maximum differences are given, based on three identical pairs within a die.

(a) All-pass rings

(b) MZIs

Figure 8.4: Within-die uniformity: width differences at short distance of each other, as a function of isolated wire width, based on the data in figure 8.3. The DC component was subtracted from the data in order to compensate for a deterministic component.

(a) All-pass rings   (b) MZIs

Figure 8.5: Within-die uniformity: differences between mean stopbands of filters at long distance of each other, as a function of isolated wire width. The minimum, median and maximum differences are given, between three identical pairs within a die. Each value is based on the mean stopband wavelength of an identical pair.

mean width changes of up to ± 15 nm. However, on a length scale of a few mm, silicon top layer thickness variations may have an influence, as noted in chapter 3. Based on the data in chapter 2, the spectral shifts can be explained fully by mean height variations of ± 6 nm. From the measurements in chapter 3, we can see that indeed over a length scale of 5 to 10 mm, variations of up to 0.6 nm can exist, explaining shifts of up to about 1 nm. As this can only explain part of the spectral variations measured, we can conclude width variations or sidewall slope variations on a longer length scale must exist and be larger than the short-range variations.

8.1.2   Non-uniformity over the wafer

A limited experiment was performed in order to quantify the non-uniformity of wires over the wafer. The channel wavelength reproducibility and channel isolation were characterised for a ring resonator, a lattice filter and a Mach-Zehnder device. The device properties are given in table 8.1. All three devices were on the same mask and fabricated on the same wafer. The structures were not covered with oxide, which may have an impact on the measured non-uniformity due to the presence of dust or chemicals. The channel wavelength and isolation were measured for three sets of nominally identical samples. The three sets were taken from different columns on the wafer and thus differed in dose used during lithography (linewidth). This is shown in figure 8.6. The different dies in a single column are fabricated with identical process parameters, and can be used
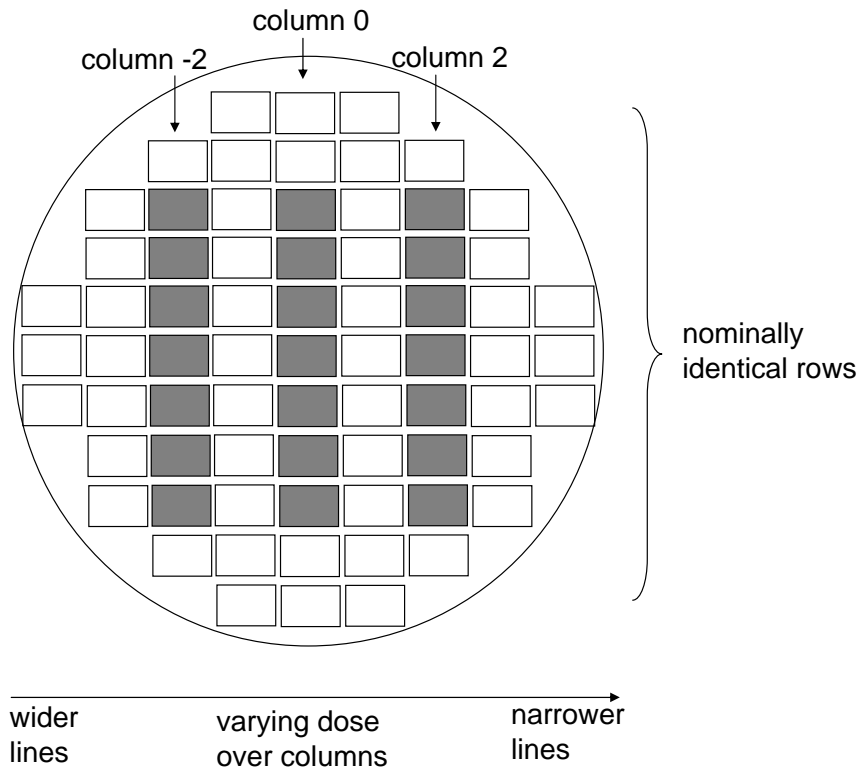
Figure 8.6: Wafer layout for characterization of uniformity within a wafer. Mach-Zehnder, ring resonator and Mach-Zehnder lattice filters were characterized on three different columns of dies.

to quantify non-uniformity over the wafer in North-South direction. By characterising different columns, the non-uniformity may be quantified as a function of linewidth.

Observed spectral variations

Figure 8.7(a) plots the spectra for an MZI with a delay length of 50 µm on the central column on the wafer. The standard deviations of the maximum and minimum transmission wavelengths are more than 2.5 nm. Note that the extinction is only 3 to 4 dB, originating from the large deviation of the directional couplers from 3 dB coupling. Two all-pass ring resonators showed larger shifts, with a standard deviation of 3 and 2.77 nm (figures 8.7(b) and 8.7(c)). The resulting standard deviation for each device measured is listed in table 8.1. Two lattice filters show a spectral shift

Device	Parameter	Column -2	Column 0	Column 2
MZI $\Delta L = 50\,\mu m$	$\sigma_{\lambda_{max}}$	3.21 nm	2.68 nm	3.47 nm
	$\sigma_{\lambda_{min}}$	2.97 nm	4.20 nm	3.74 nm
Lattice, N=11 $\Delta L = 46.1\,\mu m$	$\sigma_{\lambda_{max}}$	1.27 nm	1.26 nm	1.49 nm
Lattice, N=5 $\Delta L = 188.9\,\mu m$	$\sigma_{\lambda_{max}}$	0.7 nm	0.83 nm	
All-pass ring R=5 μm,L$_c$=4 μm gap=180 nm	$\sigma_{\lambda_{min}}$	3.2 nm	3 nm	5.32 nm
All-pass ring R=5 μm,L$_c$=4 μm gap=220 nm	$\sigma_{\lambda_{min}}$		2.77 nm	

Table 8.1: Within-wafer uniformity: standard deviation of measured band centre variations between identical structures on nominally identical samples.

with a σ of 1.26 nm and 0.83 nm. However, the filter with the smaller delay length and smaller maximum delay showed the largest variation. The FSR of the second lattice filter was very well reproduced, within the measurement error. In the two other sample columns, filters showed a very similar behaviour, with the ring resonators and the MZI showing double the variation of the lattice filter. From the figures in section 4.3.2, such a shift can be explained by a mean width variation in the 0.7 nm to 3 nm range.

Statistical analysis

The variations within a single column visually seem to be stochastic. By calculating the correlation between position and shift and checking the concordance between the wavelength shifts of the three filter types for each column, this can be analysed more quantitatively.

Correlation between position and wavelength shift In order to check for a deterministic component which may not be directly evident, the correlation between the north-south position of the die and the wavelength shift with respect to the mean wavelength was analysed for each device type (APF,MZ,LATTICE). Also, the correlation between the distance of the die from the wafer centre and the wavelength shift was calculated. As the positions are not distributed normally, and we cannot assume a normal

(a) MZI, $\Delta L = 50\,\mu m$

(b) All-pass ring, R=5 μm,L$_c$=4 μm,gap=180 nm (designed)

(c) All-pass ring, R=5 μm,L$_c$=4 μm,gap=220 nm (designed)

Figure 8.7: Within-wafer uniformity: example measured transmission spectra of identical devices on nominally identical samples.

distribution for the shifts, a rank correlation coefficient was used. The shifts were assigned a rank first, with the largest shift from the mean having the highest rank and the rest being ranked relative to the highest rank, taking into account the relative wavelength differences. Kendall's τ rank correlation coefficient was used [1]. This quantifies the possibility of a linear relationship between position and shift for each filter type. However, as the number of samples is very low, the numbers need to be taken with a large grain of salt. The results are given in table 8.2 for the correlation with the absolute position and in table 8.3 for the correlation with the distance from the wafer centre. The values in table 8.3 are rather low, or too wildly different, indicating little or no clear linear correlation. The correlation with the absolute wafer position is low or wildly differing too, except for the 11-th order lattice. However, for columns 0 and -2, the correlation for this device with the distance from the centre is equally high. So, only the correlation of 30% for column 2 holds. However, the chance that such a correlation could appear by chance is high. In conclusion, there is no clear indication for a linear correlation between wavelength shift and relative nor absolute position on the wafer. Combined with the measurement data for the silicon layer thickness deviation in chapter 3, which clearly shows a different mean thickness near the wafer rim, this learns that stochastic width or slope variations are currently larger than silicon layer thickness variations.

Concordance between shifts of all filter types   An additional test is to check to what extent the different devices experience a similar wavelength shift, in other words to quantify the concordance between the measurement sets of the different devices. A test statistic to calculate the concordance is the Kendall W coefficient, working on ranks or scores[2]. Again, for each filter type, each shift is assigned a score, where the maximum deviation from the mean has the highest score and the others are ranked relative to that score, taking the relative shift magnitude into account. The result is

---

[1]Kendall's τ is defined as $\tau = \frac{2 \cdot P}{0.5 N (N-1)} - 1$, with N the number of pairs, and P the sum over all pairs of the number of pairs with both ranks higher than the given pair. For example: if the ranks are

$$rank_a = [1, 5, 2, 3, 4, 6] \tag{8.1}$$
$$rank_b = [3, 1, 5, 2, 4, 6] \tag{8.2}$$

, then $P = 3 + 1 + 1 + 2 + 1 = 8$, and $\tau = 0.07$

[2]Kendall's W coefficient gives the concordance between sets of "observers" (in this case filter types) ranking the same "stimuli" (in this case assigning a wavelength shift to each position on the die)

	Column -2	Column 0	Column 2
MZI $\Delta L = 50\,\mu m$	0.82	0.59	0.15
Lattice, N=11 $\Delta L = 46.1\,\mu m$	0.29	0.33	0.29
Lattice, N=5 $\Delta L = 188.9\,\mu m$	0.14	0.04	
All-pass ring R=5 μm,L$_c$=4 μm gap=180 nm	-0.85	0.18	-0.33
All-pass ring R=5 μm,L$_c$=4 μm gap=220 nm		0.29	

Table 8.2: Within-wafer uniformity: Kendall correlation of the ranked measured band centres with ranked absolute position on the wafer.

	Column -2	Column 0	Column 2
MZI $\Delta L = 50\,\mu m$	-0.31	0.04	0.0
Lattice, N=11 $\Delta L = 46.1\,\mu m$	0.26	0.43	-0.05
Lattice, N=5 $\Delta L = 188.9\,\mu m$	-0.54	-0.12	
All-pass ring R=5 μm,L$_c$=4 μm gap=180 nm	-0.29	-0.13	-0.31
All-pass ring R=5 μm,L$_c$=4 μm gap=220 nm		-0.18	

Table 8.3: Within-wafer uniformity: Kendall correlation of the ranked measured band centre with ranked distance from the wafer centre.

column	W	significance
0	0.21	>0.05
2	0.11	>0.05
-2	0.08	>0.05

Table 8.4: Within-wafer uniformity: Kendall's W coefficient of concordance for the three columns measured. For each column, the spectral shifts a MZI, lattice and APF filter were compared.

given in table 8.4. The devices give a relatively significant concordance for the central column, of about 20%. For the eastern and western columns, the concordance is lower but still significant. Combined with the correlation data, this probably points towards an within-die non-uniformity that is smaller than the non-uniformity over the wafer, rather than to an actual non-stochastic dependency of the shift on the wafer position.

Conclusion

The correlation values indicate that no linear relationship between the wavelength shift and the north-south position on the wafer or distance from the centre can be seen. Given the very low correlation or highly differing values, it can be assumed we cannot see a simple non-linear relationship either. The number of samples for each filter type is too small to allow for any periodicity analysis. There is a small, but nevertheless significant concordance between the shifts of all filter types with respect to their relation to the position on the wafer.

8.2 Solving inaccuracy

8.2.1 A better fabrication technology

Several points can be tackled to improve the waveguide and structure quality.

- **Sidewall roughness.** The width variations and other roughness on a small length scale should be reduced, in order to reduce waveguide losses. For simple interconnections, the current losses are low enough, but advanced filters, such as filters with multiple coupled or uncoupled ring resonators need a lower propagation loss. This is a matter of optimising both the lithography and etching process.

Some roughness may exist on the reticle already, which is difficult to improve upon without using higher-end and thus more expensive technology. Still, it should be possible to reduce the roughness to the nanometre level, and with polishing methods it is probably possible to reduce it to sub-nanometre scale.

- **Sidewall passivation.** In order to reduce the losses further, the dangling Si bonds at the sidewall should be passivated in order to avoid absorption.

- **Longer length-scale width variations.** More important than reducing sidewall roughness is to reduce the width variations on a length scale of 100 nm and longer, as these can give rise to serious phase errors resulting in crosstalk or loss, inaccuracy between structures and non-reproducibility between samples. Reducing these could hopefully also make a significant difference for directional couplers. However, part of these variations may already be present on the reticle, and increasing the reticle accuracy comes at an increased cost. Both decreasing the amplitude and the typical correlation lengths of the width variations is needed.

- **CD variation**. Part of the CD variation between structures is due to stochastic width variations within a single line, but other effects can play that should be reduced. New process development has already shown that it should be possible, at least down to the resist level, to obtain an accuracy of the order of 1 nm over the whole wafer except for the border dies (which is within the measurement error of the SEM).

- **Sidewall verticality.** Currently, the sidewalls have a slope of 9°, which gives rise to a serious difference from what is designed, and can also explain part of the non-proper functioning of directional couplers.

- **Si layer thickness uniformity.** Obtaining SOI wafers with less roughness and a higher uniformity of the silicon top layer over the wafer could reduce the accuracy and reproducibility, as was mentioned in section 3.4.5. Currently, those variations can account for a considerable fraction of WDM channel spacings, and maximum expected spectral non-uniformity is of the order of 20 nm. Additionally, more reproducible top layer thicknesses, currently with a σ of 6 nm, could drastically improve on the wafer-to-wafer reproducibility of wavelength dependent structures. While this is of less importance in a

research environment, it is of course extremely important for real applications, as we will see it is not straightforward to compensate for such variations. However, more accurate SOI wafers comes at an increased cost.

It should be noted that all of these improvements, except for the Si layer thickness uniformity, are the subject of ongoing research.

8.2.2 Tuning

As a certain degree of phase errors is inevitable, they should be compensated for by tuning components. Such tuning is static, or slowly varying in time, and can also compensate for irreproducibility due to stochastic and deterministic processes. For complex integrated circuits, a thorough characterisation and setting of the tuning parameters is needed, which should be automated or semi-automated through an extensive calibration process, possibly combined with live, on-chip monitoring of the signals.

Temperature tuning

As discussed in chapter 2, the temperature dependence of SOI filters is high. Therefore, temperature tuning can induce a considerable phase change in relatively short length of wire, and both spectral shifts and phase errors can be compensated for. With an effective index change of around $2 \cdot 10^{-4}/K$, the phase change is of the order of $8 \cdot 10^{-4}/\mu m K$. In other words, the phase change induced by a $10\,°C$ temperature change in a $10\,\mu m$ long wire is around $\frac{\pi}{40}$.

So, even though the temperature dependency of SOI structures is high, if we want to keep the temperature change limited, a long enough wire is needed to compensate for large phase errors, of the order of some tens of micrometre. Fortunately, the waveguides in the studied lattice and AWG filters are long enough, so that only a few degrees of temperature change in the waveguide core is needed.

To check if the 70 to 80 pm/K wavelength shift is a good prediction (see section 4.3.2), the temperature-dependent shifts of AWG, Mach-Zehnder lattice and ring resonator devices were measured. The samples were put on temperature-controlled stages, so that the whole sample was heated or cooled. Some resulting measured spectra are plotted in figure 8.8. As the relative phase between the stages of a lattice filter stays constant in this way, the sidelobes in the transmission spectrum are not altered other than through a changed dispersion. For each filter, the passband or stopband

(a) Pass port of lattice filter

(b) Drop port of ring resonator

(c) Transmission of all-pass ring resonator
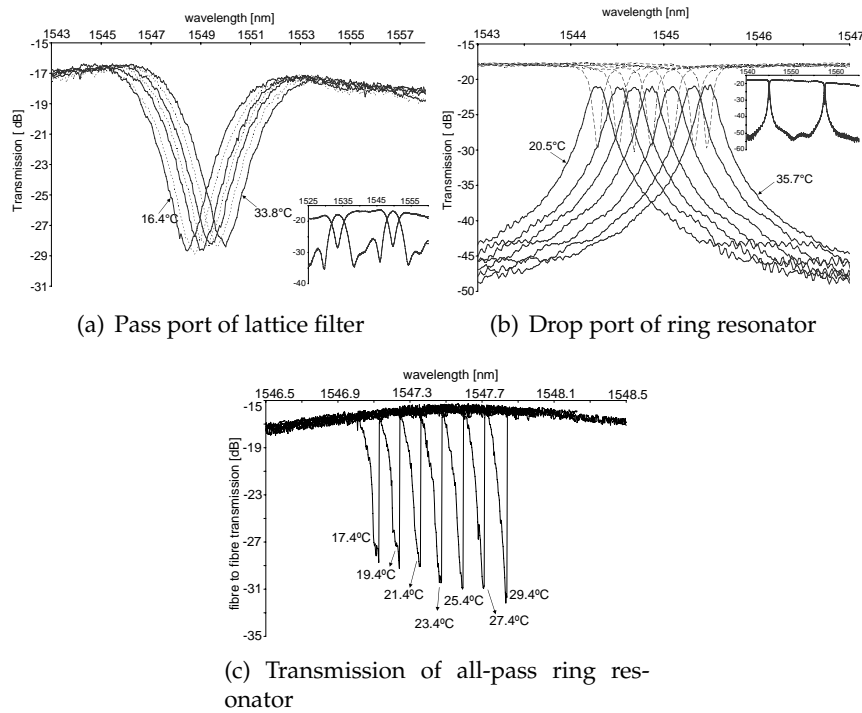
Figure 8.8: Measured filter transmission spectra as a function of temperature

centre was monitored. The result for several fabricated devices is shown in table 8.5. The measurements show a fairly good agreement with the predicted shift, with shifts between 58.6 and 89.1 pm/K. There is no clear dependency of the shift on the isolated waveguide width, but the lattice and ring filters have a coupling region where the waveguide is smaller and the waveguides are coupled, with different influence on the shift.

Tuning over the complete free spectral range of a filter is impossible, as the temperature difference would reach 100 °C. However, tuning over a WDM channel spacing of 0.1 nm (12.5 GHz) to 1.6 nm (200 GHz) is possible with a temperature shift between 1 and 20 °C. In conclusion, a nanometre scale wavelength inaccuracy can be compensated for by temperature tuning. For larger errors, another tuning method is needed, or a better fabrication technology is needed with an accuracy on waveguide dimensions to about 1 nm. As we have seen, the latter will only become possible with large process windows and with a small distribution of wire widths in the design.

device	wire width [nm]	band centre [nm]	shift [pm/K]
awg	600	1532.5	69.2 ± 3.4
lattice	560	1533	81.4±2.8
		1550	89.1±2.6
lattice	520	1541	72±0.6
AD ring	460	1552	83.4±2.2
AD ring	470	1540	85.4 ±6.2
		1551	88.2 ±1.5
AD ring	550	1545	80.4±0.6
AD ring	560	1538	79.4±3.9
		1555	82.5±3.1
AD ring, with BCB top cladding	565	1553	77.0±4.0
AP ring	548	1547	58.6±0.2

Table 8.5: Temperature dependent spectral shift of filter transmission spectra for different filter band centres and isolated wire widths.

Other tuning methods

Linear electro-optic effect   While normal silicon does not exhibit a linear electro-optic effect (Pockels effect), strained silicon can show a significant second-order non-linearity. While the SOI is strained already, the magnitude is too small to generate a large static electro-optic effect. Stress can be induced by simply applying an overcladding such as Si_3N_4 [95], an operation and material which are in principle CMOS compatible. An externally applied voltage then changes the refractive index seen by the propagating mode and thus introduces a phase change. This effect can be used as the basis for a modulator, but in principle, tuning is possible too, at least for long enough waveguides. An advantage of this approach is that only an external field is needed, and no current flows through the waveguide. A disadvantage is the need for two electrodes close to the waveguide in order to keep the needed voltage low. Two electrodes take more die real-estate than a single electrode needed for temperature tuning.

Current injection   Another option is to inject current in the waveguide, changing both absorption and the refractive index through free carrier effects. Again, this is mainly an effect used for modulators, but it can in

principle be used for tuning. However, now current is flowing through the waveguide and absorption is increased.

Overlay materials   A last attractive approach is the use of overlay materials such as liquid crystals. Liquid crystals (LC) are naturally birefringent, and the axes of *nematic* LC molecules can align under the influence of an external electric field. When applied as the cladding of an SOI waveguide, the refractive index the mode tail feels can be changed by altering the alignment of the LC. However, the molecules also align under the influence of nearby materials and geometry, such as the waveguide and waveguide edges, which makes modelling not straightforward [96]. If all molecules would properly align as desired, a large effective index difference of 0.15 would be possible for a narrow wire. However, due to the different alignment near the waveguide surface, exactly where the field is higher, the effect is smaller. Still, a colleague measured a $2 \cdot 10^{-3}$ index difference on a 3 μm wide waveguide [97], which is still a strong effect.

Of course, applying liquid crystals is not straightforward, as one needs a way to fix them to the substrate, and both the LC and the way of attaching it may hinder other structure such as electrical or fibre contacting. Other overlay materials, especially liquid ones, may suffer the same drawback. In order to apply liquid crystals very locally, microchannels and underetched structures can be used for instance, or the material can be mixed with a polymerisable material. Application over larger areas, leaving open areas for fibre coupling, is easier and is already used in Liquid Cristal on Silicon (LCOS) technology.

Current problems that cannot be solved through tuning

Some problems are difficult or practically impossible to compensate for by tuning mechanisms. Therefore, the approach of just putting heaters on everything that does not seem to work is not an ideal solution, and not even a practical solution for many structures. In such cases, better fabrication technology, in combination with better design, is the only option.

- **Imperfect directional coupling** While there is an intrinsic problem to exciting a high-index contrast directional coupler, definitely for small gaps and small bend radii, it is not expected to be a main problem at least for reasonably wide gaps. However, the fabrication variations make control over the coupling difficult. A fabrication process with larger process windows and more vertical sidewalls, should improve a lot on the directional couplers.

- **Backscattering in resonators** In chapter 5, it was shown that even very limited distributed backscattering at sidewall roughness can lead to serious problems for a large range of ring resonator configurations. This should be reduced by better fabrication technology.

- **Large spectral shifts** In this chapter, we have seen that shifts of the order of ± 10 nm can exist between components spaced millimetres apart. It becomes difficult to overcome this through active tuning. For instance with temperature tuning, it is possible, within a typical temperature budget, to compensate up to a few nm. Therefore, the variations should be reduced to the O(1 nm) level, so that residual variations can easily be compensated for. However, silicon layer thickness variations are beyond our control, and thus pose a fundamental problem, if a large part of the variations over a longer length scale can be attributed to them, as mentioned above.

8.2.3   Intelligent design

In addition to fabrication process improvement, one can take considerations on non-uniformity already into account in the component and mask designs. For instance, waveguide cross-sections can be changed or high-level parameters can be adapted. But various trade-offs will play in doing so, and therefore for actual devices such design changes cannot be done blindly. Increasing tolerance can either be at the core of the design procedure, or only loosely related to it.

An example of a more low-level design option for greater tolerance is the broadening of wires in lattice filters (chapter 6) and AWG devices (chapter 7). This has a small influence on the high-level component design, as it changes the effective index and group index of the waveguides of the delay lines, and thus it needs to be taken into account when designing for a given FSR or passband/stopband wavelength. Also, it has an influence on the low-level implementation of the component, as tapers need to be added, putting limits on the achievable component size, configurations and parameters. As already mentioned in section 6.2.4, with CMOS fabrication processes, it is difficult to get both the narrow and wide waveguides exactly on target. Therefore, control over the exact wavelengths or FSR, or control over for instance coupling coefficients, is traded for a phase error reduction. Another example is increasing the coupling gap in a directional coupler in order to make the width and gap more controllable. Even with a higher resolution lithography process, this may be needed in view of the variety of structures ending up on a single litho layer.

On a higher level, filters can be designed for a greater tolerance against variations. For instance, in a lattice filter, it is possible to design the filter coefficients so that the resulting filter has a larger tolerance against coupler variations. Also, additional stages can be added to compensate for a high crosstalk, although there is an optimum as seen in chapter 6. A similar approach could be used for ring lattice filters. In an AWG, less high-level design options are possible.

Another option is to rely on things than can be very well controlled, for providing the functionality of the device. For instance, a major advantage of optical projection lithography compared to e-beam or ion beam lithography is the great control over distances, such as the pitch of a one-dimensional or two-dimensional grating. Therefore, implementing a filter function using differences between pitches may be a very viable option.

Yet another possibility is to rely on the difference between feature sizes. Even though it is difficult to have good control over different feature sizes at the same time, it could be interesting to see whether devices relying on differences between feature sizes are more tolerant than devices relying on the exact feature size of a single feature as such.

8.2.4 Conclusion

Putting the data from chapters 2, 3, 6, 7 and this chapter together, some conclusions can be made.

The different data strongly point in the direction of stochastic linewidth variations along a wire with a magnitude of about 5 nm (in terms of width), with a typical correlation length along a sidewall of 500 nm to 1 μm. This is supported by

- Top-down SEM pictures of straight wires (Ch. 3)

- The correlation length and magnitudes reported in [52] on line edge roughness for very similar 4X masks (Ch. 3)

- The magnitude of the CD variations measured by the mask shop on our own reticles (Ch. 3)

- The magnitude and variance of the phase deviations in Mach-Zehnder lattice filters and AWG devices (Ch. 6 and 7).

- The non-uniformity over a short distance, measured in this chapter, which can be explained by mean linewidth variations, which can be related to such a stochastic width variation.

If the source of these variations is already on the mask, as suggested, the solution is beyond our control. It is possible to switch to a mask accuracy of $\pm 0.014\,\mu m$ (on reticle) instead of the current ± 0.022 μm, at relatively limited cost, depending of course on the quotations given by the mask shop. Whether this would actually reduce the supposed variations is at this point unclear. The impact on the short-distance-range reproducibility would be limited, but the phase errors in filter devices could be reduced. Instead of $0.18\,\mu m$ CMOS technology, we can also switch to $0.15\,\mu m$ or $0.13\,\mu m$ technology, with better CD uniformity, but at a high additional cost.

The reproducibility of identical components depends on the distance between the components, with a near-scale reproducibility on a distance of a few tens of micrometre that is clearly smaller than that over a length scale of a few mm. These are due to stochastic mean CD variations or sidewall angle variations, for instance due to local illumination dose or resist variations, and to silicon layer thickness variations. From the thickness measurement in chapter 3 and the data in chapters 2 and 4, it can be concluded that height variations can only account for differences over a few mm distance on the nm level, and width or slope variations are larger on a longer length scale. This is supported by the die-to-die-variations, where no clear dependency can be seen on the position on the wafer.

As silicon thickness variations are beyond our control and width variations currently seem to be stronger, processing accuracy should be improved to the point that all mean width variations are of the order of just a mere few nm, so that tuning methods can be used effectively to compensate for them.

Chapter 9

Applications

Besides for telecommunication, SOI wavelength filters can be applied in other fields, and have different performance requirements that could be met sooner than those of the telecom world. In this chapter, such applications are studied. In particular, two examples are given. First, a polarisation independent AWG-based duplexer for use in access networks is discussed. A more extensive discussion of an AWG based router for high-speed interconnect networks follows. Both devices use WDM technology (channel spacings, wavelengths) for communication over shorter distances.

The duplexer was designed and fabricated using the simulation code discussed in chapter 7, and for the first time, a functional polarisation independent component in SOI is demonstrated, based on polarisation diversity.

The AWG router was fabricated, fibre pigtailed, characterised and used in the context of an ESA project researching the use of optical backplanes for board-to-board communication in spacecraft at multigigabit speed. While the use of SOI technology in space is possibly still far off, this project formed a nice framework for studying the possible performance in a real-world application. Furthermore, an initial study of the radiation hardness of SOI components was done. This is discussed in section 9.3.

9.1 A polarisation independent duplexer for access networks

9.1.1 A WDM system for fibre-to-the-home

In fibre-to-the-home systems currently deployed over a passive optical network (PON), upstream and downstream traffic are typically modu-

239

Figure 9.1: A possible duplexer based access point for FTTH applications. The downstream and upstream data are each on a channel on the 50 GHz ITU grid. Eight wavelengths are contained within broader downstream and upstream bands, to support multiple users over a single fibre. Wavelength allocation is done at the central office and modulation is done at the access point.

lated on two or three separate wavelengths, with downstream traffic in the 1550 nm band and upstream traffic around 1310 nm. ATM (aPON) or Ethernet (ePON) are then used for access control. However, by using a WDM technique over the PON (wPON), where multiple wavelengths are used for multiple users over the same fibre, data capacity can be enhanced and flexibility enlarged. An elegant setup is when all users have the same, wavelength-independent transceiver, and wavelength allocation is done at the central office. The upstream data can for instance be modulated on a downstream CW carrier by a reflective SOA [98]. This is illustrated in figure 9.1. A possible scheme is to allocate densely spaced channels within an upstream and a downstream band. With the modulator working wavelength-independently within the upstream band, it does not need to know which of the specific wavelengths it is acting on. A duplexer is needed to separate upstream and downstream traffic.

9.1.2   Design of an AWG-based duplexer

Such a duplexer was designed, with two 400 GHz wide bands, separated by 150 GHz. Each band can contain eight 50 GHz channels, which the duplexer is unaware of. In this way, it is for instance possible to scale the system to 16 channels with 25 GHz separation by only modifying the central office hardware. However, designing a filter with two fairly wide bands with small spacing is not straightforward. Therefore, an AWG was used with two channels used to compose a single band. At the output of the AWG, both channels can be combined into the upstream or downstream
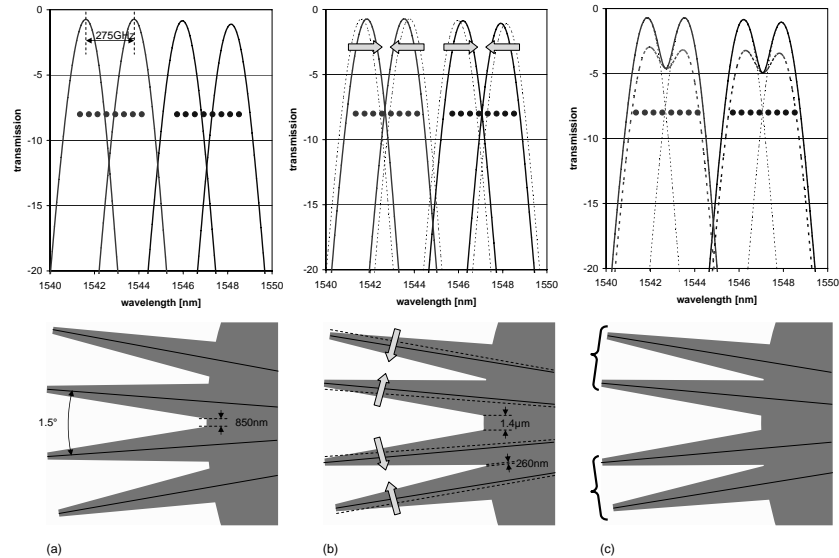
Figure 9.2: An AWG-based duplexer: optimisation of the port locations to obtain two 400 GHz bands, spaced by 150 GHz. (a) Original AWG port locations and the four corresponding transmission bands. (b) The band pairs are moved closer together by moving the ports. (c) The port pairs are summed to obtain the upstream and downstream bands.

band by combining them electrically or through a 2-D fibre coupler, avoiding interference.

The AWG was designed with a nominal channel spacing of 275 GHz, and a FSR of 3.85 THz. The input and output ports of the AWG were displaced so that the two channel pairs combine into the two 400 GHz bands needed. This is illustrated in figure 9.2. The AWG has 48 arrayed waveguides and has a footprint of 600 μm × 350 μm. The double etch scheme was used to reduce insertion loss, and the arrayed waveguides were broadened to reduce crosstalk (see chapter 7). The, design of the AWG was done by simulation of the transmission spectrum, in a try-and-repeat process. The component was designed by Dr. W. Bogaerts, using the computer code developed in this work (chapter 7).

For FTTH systems, polarisation independence is very important. As the waveguides and circuits are inherently polarisation dependent, a polarisation diversity approach was used. The two orthogonal polarisations in the fibre are coupled to separate waveguides by a 2-D grating fibre coupler [39]. From the point of view of the waveguides, both polarisations are

(a) Schematic

(b) SEM

Figure 9.3: An AWG-based duplexer, with 2-D fibre couplers for polarisation independent behaviour. The two polarisations are propagated through a single AWG in opposite directions.

identical, and a polarisation independent circuit can be realised by looping both paths through identical components. At the output, both paths can be combined by using an identical 2-D fibre coupler. As both paths stay orthogonal, no interference results. In the case of this single-component AWG duplexer, an elegant approach is to propagate both paths through the single AWG, in opposite direction. This avoids the need for a second, identical AWG, reducing footprint and avoiding polarisation dependency due to fabrication variations. With separate AWGs, sub-nm control over waveguide widths and heights would be needed, or tuning of the two devices, in order to achieve polarisation independence. With a single AWG, only polarisation dependent loss (PDL) can result, due to the extinction ratios of the fibre couplers combined with possible loss variations in the input and output waveguide paths. The scheme is illustrated in figure 9.3.

(a) Transmission spectra (b) Polarisation dependence
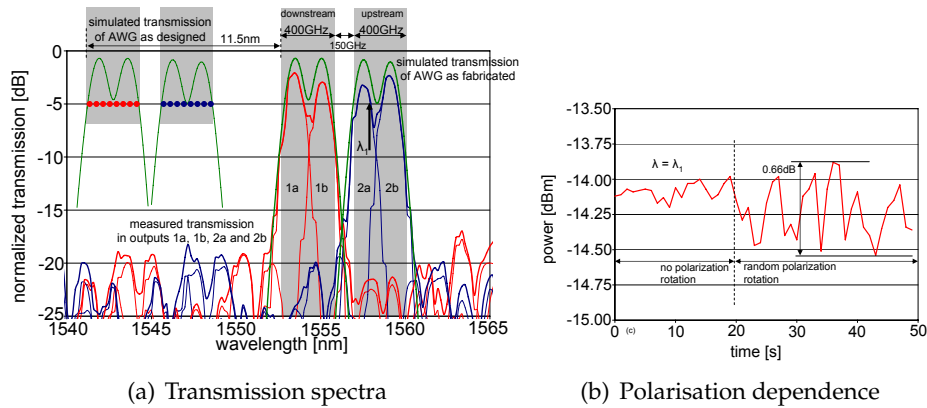
Figure 9.4: (left) Combined normalised transmission spectra from input to the four output ports of the duplexer. The bold line shows the summed transmission, obtaining the downstream and the upstream bands. (right) Power as a function of time, for randomly varying polarisation.

9.1.3 Experimental result

Figure 9.4 shows the combined normalised transmission spectra measured through the two output pairs, and the simulated transmission. When the measured output powers for each pair are added up, simulating electrical combination or combination through a fibre coupler, the non-uniformity within a band is 3.4 dB, with an insertion loss between 2.2 dB and 5.6 dB. The isolation between upstream and downstream bands reaches 15 dB. The fibre coupling loss is 6.7 dB, with a 60 nm 3 dB bandwidth. The absolute positioning of the bands is not well controlled yet, however. This can be compensated for by tuning. As expected, no polarisation dependent spectral shift could be measured.

To measure the PDL, the output power was monitored while randomly varying the input polarisation. Figure 9.4(b) shows the power fluctuation over time with and without polarisation rotation. The wavelength was kept constant near the peak of port 2a. Index-matching fluid was applied for stabilisation of the input and output fibres. The measured PDL is only 0.66 dB.

9.1.4 Conclusion

By use of 2-D grating couplers, propagation of two light paths in opposite direction through an AWG, and optimisation of the input and output ports

of the AWG through simulation, a polarisation independent duplexer for use in FTTH applications was designed. The device was fabricated using CMOS-based processes including deep UV lithography. While the wavelength inaccuracy is still of the order of some nm, the performance of the device matches simulation well, and no polarisation dependent wavelength shift could be seen. The polarisation dependent loss is 0.66 dB.

9.2 Router for a multigigabit optical backplane

Optical datacommunication between electronic boards can greatly enhance the data transfer capacity compared to electronic links. While the bandwidth of electronic links is pushed further and further, they are limited to short distances. The bandwidth × distance product of optical links is much larger, putting virtually no limits on the distance between boards. For optical datacommunication, typically VCSEL sources operating around 980 nm are used in combination with multimode fibre. The use of multimode fibre is usually justified by the better alignment tolerance compared to single-mode fibre. Parallel links by using VCSEL arrays have been explored [99]. However, these systems are often point-to-point or use a broadcast-and-select architecture. The application of WDM for the creating of an interconnect network with wavelength selective routing has been examined. While VCSEL sources can also be used in a coarse WDM approach, the use of telecommunication technology is an interesting option that can possible offer higher data aggregate data rates and greater flexibility.

While optical board-to-board technology is not yet extensively used, it can be expected to emerge at some point in the future. For example, the requirement for data handling on board of spacecraft is expected to increase dramatically in the near future. A first development driving this evolution is the introduction of telecommunication satellites with a larger number of (finer grained) channels or channels with larger data capacities. The power consumption of the (electrical) on-board processors routing the data is growing extremely high and is dominated by interconnect losses. A second evolution is the implementation of high-data-rate cameras, radars and other sensor devices. More and more, the acquired data will be processed on board before communicating. At some point, the power, weight and size of an optical communication interconnect network will be (much) lower than for an electronic network offering the same functionality.

In the ESA research project *Multigigabit optical backplanes*, the implementation and use of a wavelength-routed high-speed optical board-to-

Figure 9.5: Routing function of the backplane. Each node has a sender (left) and a receiver (right). A cyclic AWG is used, so that N wavelengths can be used to make N concurrent but mutually exclusive data connections. This scheme can be expanded to include broadcasting or multicasting communication channels.

board interconnect backplane was studied. A demonstrator was built, using an SOI AWG wavelength router as the core of the backplane. An SOI device was fibre pigtailed and the demonstrator was characterised using bit error rate measurements, giving an indication of the scalability of SOI technology at the current level. Also, first measurements were performed in order to test the radiation hardness of SOI components. An important problem in using SOI for space applications is the polarisation dependency. This problem was not tackled here, although polarisation maintaining fibres could be used for this application.

9.2.1 Backplane architecture

The optical backplane connects a number of nodes (boards) through a central routing device. Through the use of wavelength addressed routing, the optical power needed is minimised compared to power splitting networks. Each node is equipped with a tunable laser and a (fixed) receiver, each connected through a single-mode fibre to the wavelength router. With this setup and a cyclic router, the number of wavelength channels needed is kept O(N) for unicast communication. The scheme is illustrated in Figure 9.5. Multicast or broadcast communication can be provided for at the expense of a higher number of wavelength channels and a higher power loss due to the introduction of wavelength-independent splitters and combiners.

9.2.2 Demonstrator

The demonstrator consisted of four nodes, two of which were equipped with a widely tunable laser and a detector. Two boards only had receiving

capability. The demonstrator setup is schematically shown in Figure 9.6. The backplane streams data at 10 Gbps (uniplex) using 10 Gbps modulators. Each node consist of an electrical board equipped with an FPGA, microcontroller, switches and many electrical interfaces, such as the high-speed MSA bus for connecting to the transponders. The boards fulfil a double function implementing both the control plane and data plane of the network. For the demonstrator, a simple hybrid centralised/distributed traffic control mechanism is employed, where one node acts as a master triggering the other boards but each board contains its routing table, controls its laser, modulator and receiver and generates, receives and handles data. The central triggering is important in order to avoid synchronisation problems. The boards generate 10 Gbps data, send it out over the right wavelength, and receive data at 10 Gbps. Both control and data plane functions are implemented in the FPGA. All boards generate the same data stream consisting of a preamble and pseudo random data generated by a random generator implemented in each FPGA. The preamble enables the receiver and control electronics to lock in on the signal. The received data is compared to the data generated by a second random generator (using the same seed) and the total number of sent bits and number of faulty received bits are recorded in registers. In this way a direct measurement of the bit error rate is obtained. Internally, the boards run at 155.52 MHz with a 64-bit data path, resulting in an aggregate (but single) output stream at 9.953 Gbps. A real-world communication protocol was not developed at this level. Rather, the pseudo-random stream is generated as long as the network is not switched to another connection topology.

The tunable lasers were widely tunable CW lasers with control circuitry developed by Intune Technologies[1], with 50 GHz channels between 1530 nm and 1560 nm. For two boards, modulator, detector and receiver were integrated in a laserless transponder bought from Avanex[2]. The modulated output power is about 3 dBm. The receivers are pre-amplified PIN photodiodes with a specified sensitivity of -18 dBm, but a measured sensitivity that is about 10 dB better.

The two boards with receiving capability only had burst-mode receivers developed by Intune Technologies. Three AWG routing devices were used: a commercial grade 4 × 4 Silica-on-Silicon (SoS) device, a prototype 4 × 4 SOI device developed in this work and a polymer 8 × 8 device obtained from Heinrich Hertz Institute Berlin. The SoS device has the lowest in-

---

[1] Intune Technologies, Dublin, Ireland. http://www.intune-technologies.com. Was a partner in the ESA Multigigabit project.

[2] Avanex,Inc., http://www.avanex.com
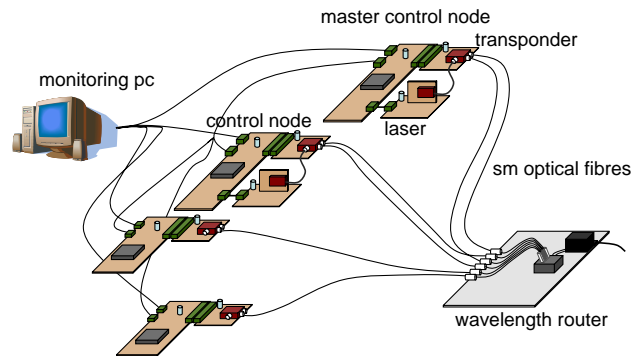
Figure 9.6: Schematic of the demonstrator. Each node consists of an electronic board with an FPGA, microcontroller, and various electrical interfaces. Each node is equipped with a laserless transponder. Two of the nodes are equipped with CW lasers. The transponders are connected to the wavelength router (AWG) through standard single mode fibre. Network control is hybrid centralized/distributed, implemented in each FPGA but with a single node acting as master. A PC is used for monitoring the network.

sertion loss and lowest crosstalk and was used as the primary device for the demonstrator. The SOI device was used to test the possibilities of high index contrast SOI technology and has the highest insertion loss and crosstalk. The polymer device has more channels and was used to test broadcast and multicast communication paths. The polymer device has an insertion loss and crosstalk in between those of the SoS and SOI components.

9.2.3   AWG

We used a second-generation AWG with shallow star couplers (see chapter 7). The device has 4 channels with 250 GHz channel spacing and was designed to be cyclic with a 1 THz free spectral range. The footprint of the AWG is 425 μm × 155 μm (0.066 mm^2), using waveguide bends with 3 μm radius. The device size is limited by the delay length needed in order to obtain the 1 THz FSR (71.36 μm) and the number of waveguides used (10). The main parameters of the AWG are given in Table 9.1

9.2.4   AWG pigtailing

In order to couple light from and to fibre, shallow grating couplers were used. Light is coupled from standard single mode fibre (9 μm core) mounted under an 9° angle to a 10 μm broad waveguide, which is then linearly ta-

Channel spacing	250 GHz
# channels	4
FSR	1 Thz
Delay length	71.36 μm
Grating order	113
Focal length	21.83 μm
Waveguide width	500 nm
Port aperture spacing	220 nm
Arm aperture spacing	190 nm
Aperture width	1.5 μm

Table 9.1: Design parameters of the demonstrator AWG

pered down towards a 500 nm wide wire. The grating couplers are designed for optimum transmission around 1545 nm, corresponding with the centre wavelength of the tunable lasers, and have a pitch of 625 nm, 20 periods and an etching depth of 70 nm. With a 750 nm top oxide cladding and an index matching material between fibre and coupler, the calculated maximum efficiency is 46% with a 1 dB bandwidth of 46 nm. The great advantages of this approach are the relaxed alignment tolerances and the possibility of wafer-scale testing.

For fibre pigtailing the device, a 1-by-8 array connector with standard single mode fibres was used during testing, and was then aligned and glued to the chip. The connector consists of a silicon V-groove assembly with a pyrex cover and fibre facets polished under an 8° angle and with a standard 250 μm fibre spacing, commercially available from OZ Optics[3]. The surface of the connector has approximately the same size as the die and one obtains an attractive approach to packaging SOI chips, as the resulting component is basically a small rigid connector with the more delicate SOI chip attached to it. This is illustrated in Figure 9.7, showing also the experimental device, which was however pigtailed to a chip cleaved much larger than necessary.

The connector is mounted under a 9° angle, which results in a 1° angle between the connector facet and the chip. In this way, the edge of the connector can safely rest on an empty edge of the chip without harming the fibre facets, and application of glue is easier. A first coarse alignment of the connector to the chip is done passively. In-plane, the connector edges are aligned (both by translation and rotation) to markers on the chip, il-

---

[3]OZ Optics, www.ozoptics.com

Figure 9.7: SOI chip with AWG pigtailed to an 8 fibre array connector. (a) Schematic. (b) Picture, where the chip is cleaved much larger than necessary.
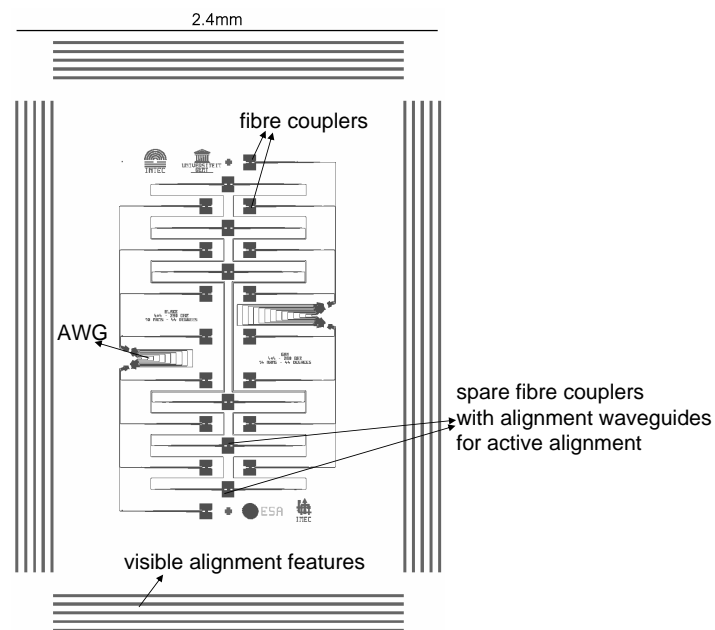
Figure 9.8: Mask layout for pigtailing to a 8 fibre array. Two AWGs are each connected to 8 fibre couplers. A spare set of fibre couplers is used for active alignment. These spare couplers are interconnected by simple waveguides two by two. At the borders, sets of lines are used for passive alignment of the connector to the chip.

lustrated in Figure 9.8. Vertically, the translation and rotational offsets are corrected through a microscope view. The passive alignment aligns the connector to a set of spare fibre couplers, interconnected two by two. Further alignment is then active, by simple transmission measurements through fibre pairs and the spare interconnected couplers. First, a coarse active alignment is done. Due to the large alignment tolerances, this is an easy and quick step. Then, the connector is displaced to the fibre couplers of the actual AWG and alignment is coarsely optimised again using a broadband light source and OSA or a laser on a channel wavelength. Next, the connector is pulled back from the chip in order to widen the diffracted spot size and relax the alignment tolerances while applying glue. The glue used in the test was a UV curable glue with a low shrinkage. UV curing is possible by illuminating through the pyrex cover over the fibres or under a certain angle. The glue also acts as an index matching material, enhancing the transmission from fibre to waveguide. A short pre-curing step forms a limited amount of crosslink bonds in the glue. Then,the connector is pushed back closer to the chip and an active optimisation of the alignment is done. Finally, the main curing step fully crosslinks the glue. Partly due to the pre-cure step, the shrinkage is limited and alignment survives the main curing step.

The connector surface size is 3.8 mm $\times$ 2.03 mm, as large as the die size with alignment structures. The die without alignment structures is 2 mm^2 in size and is only sparsely filled. In the first test performed, the chip was cleaved much larger than the die size however, resulting in a component less appealing than described earlier. However, it should be perfectly possible to dice the chip to the die size and obtain a more streamlined fibre pigtailed device. For usage in the demonstrator, the connector was put on a Peltier element. The component was thus temperature controlled through the Aluminium holder of the connector, the silicon bottom of the connector and the silica top cladding of the device. However, this thermal path was good enough for keeping the temperature stable around room temperature.

9.2.5 Performance of the pigtailed AWG router

The pigtailed AWG was characterised through static transmission measurements with an the ASE of an EDFA as broadband light source and a spectrum analyser, resulting in the spectra shown in Figure 9.9, measured two hours after gluing.

The channel spacing corresponds very well to the designed 250 GHz. The free spectral range is one tenth of a channel spacing larger than de-

signed (1.025 THz), making the AWG not perfectly circular. The total (fibre-to-fibre) insertion loss of the device is 12.5 dB for the best channels. Being a routing device, the insertion loss of the outer input and output ports (1,4,5,8) is intrinsically higher than that of the inner ports (2,3,6,7). However, the transmission spectra and the low waveguide propagation loss suggest a large part of the fibre-to-fibre insertion loss non-uniformity is due to a residual in-plane rotational misalignment of the connector to the chip.

The shortest alignment waveguide on the sample is 2.4 mm long. On a nominally identical sample, the fibre-to-fibre transmission through this waveguide (and the to fibre couplers) was -13 dB before gluing the connector to the chip. With about 0.5 dB of waveguide losses, the maximal fibre coupler transmission efficiency is then about 24% (with the air gap). The actual AWG insertion losses for the inner and outer input waveguide to all output waveguides can be calculated from transmission measurements on the AWG (taken again before gluing). One has to compare with the transmission of the alignment waveguides. However, the uncertainty on this comparison is high due to the possible difference in alignment. By comparing with the shortest alignment waveguide, and compensating for the difference in alignment and access waveguide lengths, we can calculate the upper limit for the the actual AWG insertion loss of the best channel to be 3.5 dB.(It is possibly better but difficult to measure). After addition of the (index matching) glue, the fibre coupler efficiency is expected to be higher. However, during the gluing alignment can get slightly disturbed. Still, with the best fibre-to-fibre insertion loss of 12.5 dB and an upper limit of the AWG device insertion losses of 3.5 dB, this gives an upper limit of the fibre coupler efficiency of 35% with the glue.

The sidelobe level of the measured transmission spectra (crosstalk) is still -12 dB for this AWG, mainly due to phase errors in the arrayed waveguides.

Relaxation

After 24 hours, 3 days and two months, new transmission measurements were taken in order to look for drift due to relaxation. A comparison of the spectra taken 2 hours and 2 months after gluing is shown in Figure 9.10. Slight differences can be seen and can be attributed to a slight relaxation of the assembly. However, a difference in fibre patch cable connector loss may also explain part of the insertion loss difference. The measurements were both taken around room temperature, but were not temperature controlled, explaining small spectral differences.
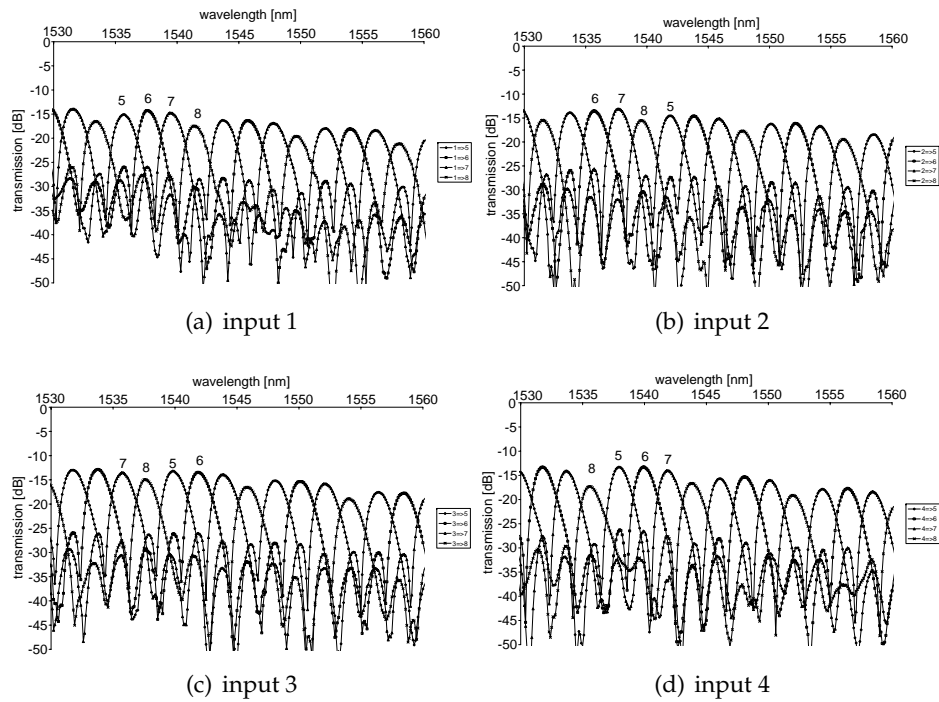
Figure 9.9: Fibre to fibre transmission spectra of the pigtailed AWG from all four inputs to the four outputs. Two hours after gluing


Figure 9.10: Comparison of transmission spectra 2 hours and 2 months after gluing

9.2.6 Backplane performance with SOI router

The backplane was tested through direct bit error measurements by the FPGAs on the test boards. The measurement setup is shown in Fig. 9.11. The BER of one communication path is monitored as a function of power incident on the receiver, once without and once with a crosstalk generating signal. The resulting single-channel crosstalk power penalty gives a first idea of the scalability of the SOI component.

Two fully equipped boards are used and polarisation controllers are used for optimal transmission. The monitored signal results either from a board communicating to itself or to the other board. An attenuator was used at the receiver side to obtain the BER as a function of received power. Due to the high receiver sensitivity, an extra attenuation of 3 to 15 dB was needed, depending on the channel combination and BER, in order to obtain measurably high bit error rates in the 10^{-7} to 10^{-12} range.

In order to measure the influence of the crosstalk level, the monitored input signal was attenuated without attenuating the crosstalk generating signal. Due to the stochastic nature and level of the crosstalk, the SOI AWG had widely varying crosstalk figures, depending on the input/output port combination[4]. Several input/output combinations of the AWG were tested, with cases resulting in coherent or incoherent crosstalk. Figure 9.12 shows some of the measurement results, with incoherent or coherent crosstalk and different (non-attenuated) crosstalk levels. Summarising all measurements, for incoherent crosstalk between -19 and -11.1 dB, a power penalty between 0.1 and 1.0 dB was measured at a BER of 10^{-10}. With coherent crosstalk, the power penalty is larger, ranging from 0.2 to 2.0 dB for a crosstalk between -20.1 and -14.1 dB.

With the SOI AWG component, only single-channel crosstalk was tested. In order to get an idea on the scalability of the system, the demonstrator was tested with an extra crosstalk contribution, however with a silica-on-silicon AWG and burst mode receivers that could be used in a more typical fast switchable backplane. These receivers had a sensitivity of -17.3 dBm at BER=10^{-10}. Based on measurements, a worst-case scenario model was proposed whereby every subsequent port in a high-port count AWG contributes an additive penalty that is equal to that of the first port, given additive crosstalk [100]. Within the power budget, a 12dB power penalty is possible. However, the above assumption of a uniform increase of power penalty with additional additive crosstalk is deemed unreliable beyond a 5dB penalty. Therefore, this value was taken as a limit to calculate the scal-

---

[4]This is in contrast to most AWG devices with decent crosstalk performance, showing a nearest neighbour crosstalk that is higher than that originating from the other channels.
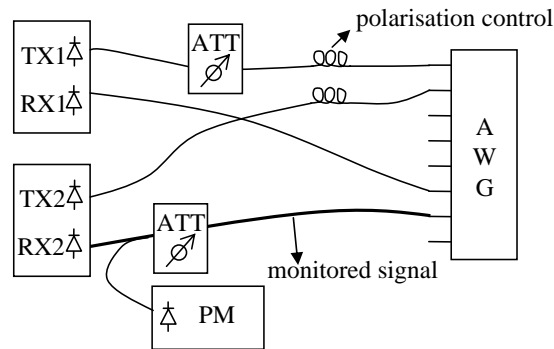
Figure 9.11: Setup for measuring the BER in the high-speed interconnect, in the presence of crosstalk. The monitored signal is the signal incident on RX2. The transmitter of the monitored signal is either TX1 or TX2, while the other transmitter generates crosstalk onto the monitored path. The power of the signal or crosstalk generating input can be adjusted, as well as the received power, in order to obtain the BER as a function of incident power on the receiver, for an adjustable crosstalk level.

ability of the technology. With the SOI component having a penalty in the 0.1 to 2 dB range, the scalability of the system with SOI technology would be just about 4 channels at a BER of 10^{-10} in a worst case scenario. Given a typical penalty of 0.5 to 0.7 dB, a scalability of 9 ports could be attainable. As the specs for the modulated output power and receivers sensitivity are considered to be representative for a typical system for this application, the second generation SOI AWG system can thus be considered to be scalable to about 4 to 9 ports. With better crosstalk values of -20 to -25 dB, scalability would rise to the 20 to 100 port range and probably better [100].

9.3   Radiation hardness testing

The sensitivity of devices to impact by radiation is a major concern for space applications. High energy radiation impacting with spacecraft originates from particles trapped in belts in the Earth's magnetosphere, from cosmic rays, from solar flares and from the solar wind. Particles trapped in belts in the Earth's magnetic field mainly include electrons up to a few MeV and protons up to hundreds of MeV. Galactic cosmic rays and accelerated solar material generate a constant, low flux of mainly protons and smaller amounts of α particles and heavier nuclei with high energies (up to several GeV). Solar flares send out γ radiation across the entire electro-

(a) incoherent crosstalk

(b) incoherent crosstalk

(c) coherent crosstalk

(d) coherent crosstalk
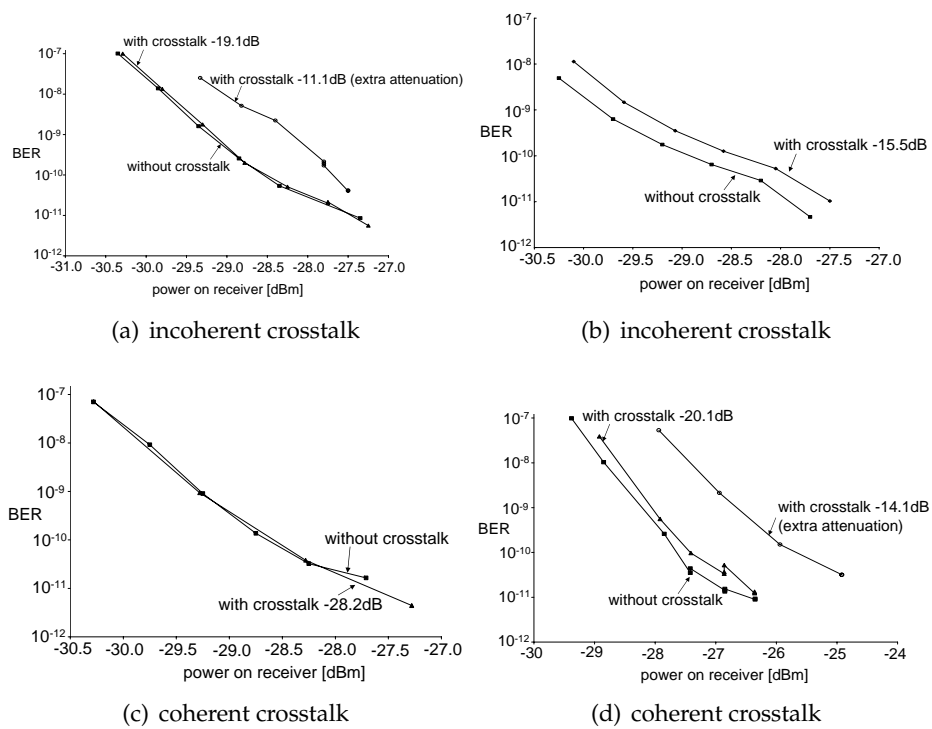
Figure 9.12: Measured BER versus power incident on receiver for 4 different input/output combinations. The BER was measured with the SOI AWG, using the setup in figure 9.11.

magnetic spectrum, often accompanied by highly energetic particles from coronal mass ejections. The solar wind plasma pervades the geo-space with a high flux of electrons and protons with energies around 100 keV. Additionally, γ radiation caused by the The slowing down of β particles (Bremsstrahlung).

9.3.1 Effects of radiation

The main optical effects we are concerned about are absorption, refractive index change and physical dimension (and density) change.

These are caused by three types of physical effects in their turn all generated by a number of different processes and metastable intermediate states, and resulting from a multitude of radiation sources (α,β and γ) [101]:

- **Total ionising dose effects.** Incident radiation with a high enough energy can break atomic bonds and create electron-hole pairs. In insulators, the charges recombine, are swept away or can get trapped in the insulator or at the insulator-semiconductor interface. The generated electrons in SiO_2 have a high mobility [102] and are swept out of the oxide very quickly, but the holes are transported much more slowly and can be trapped. In thermally grown oxide, most of the hole traps reside near the Si/SiO_2 interface. They can however be removed by tunnelling of electrons from the Si or thermal annealing. Remaining traps can have a noticeable effect on refractive index and absorption. As the SOI waveguides studied here are covered by a thin natural thermal oxide layer, absorption at sidewalls due to trapped charge may be enlarged. Experiments with glasses have shown that gamma radiation can be considered to generate a worst case effect in bulk glasses [103].

- **Displacement damage effects.** Incident particles or secondary electrons originating from γ radiation can displace atoms and create long-lasting defects. This will be the main contributing effect in semiconductor lattices. Results of experiments with 1MeV γ radiation can be explained by creation of a homogeneous density of isolated defects in the material [104]. Depending on energy and material, the created defects can potentially be rather mobile after generation, and can gather together or form dopant/vacancy complexes in doped lattices. Studies have shown neutron irradiation typically generates these kind of defect clusters. The defects will have an effect on the optical absorption and (through the Kramers-Kronig re-

lations) on refractive index. Also, the material density may change, resulting in an additional refractive index change and dimensional change. The influence of γ and neutron irradiation on the electrical behaviour has been shown to be very different. On the optical side, this is yet unclear.

- **Single-event effects.** A high energy particle passing through the SOI material may cause high local charge densities, temporally resulting in optical absorption and refractive index change, but may also introduce lasting displacement damage.

The total ionising dose (TID) received in orbit can vary greatly, depending on the altitude and longitude (for GEO missions), shielding, and random events. LEO missions with a 5 to 10 year duration can expected dose levels of 1 to 100 krad[5] when shielded with a few mm of Al. GEO missions can expect up to 500 krad over a 15-year lifetime in a worst-case scenario, but will normally absorb much less [103].

9.3.2   Experiment

In order to get a first idea of the sensitivity of nanophotonic SOI devices, a total ionising dose experiment was performed, looking for spectral shifts due to refractive index change and excess loss due to absorption. We only worked on the phenomenological level and did not try to separate different physical contributions to the effects measured. As the waveguide core is a semiconductor lattice, displacement damage effects may form an important contribution. Therefore, one needs to keep in mind that the source of effects seen in the TID experiment performed will have a large displacement damage component. Also, a test with neutron irradiation might be worthwhile as the effect may significantly different, but can be hindered by the activity of the silicon after the test. As the work here has a preliminary character, a γ TID test was deemed appropriate.

Samples were exposed by a Co60 γ source with photon energies at 1.173 MeV and 1.332 MeV and a dose rate of up to a few tens of rad/min. While this is a high dose rate compared to many situations in LEO and GEO, it is a relatively low dose rate for experiments. From these experiments we cannot know whether dose rate has a large influence on our structures. Samples were exposed up to a total absorbed dose of 300 krad (Si). For most orbital missions, this dose is much higher than the expected

---

[5]The rad is a unit of absorbed dose, still frequently used but actually superseded by the SI unit Gray. 1 rad = 0.01 Joule absorbed per kilogram of material = 0.01 Gray

**Figure 9.13:** Radiation testing setup at ESTEC, Noordwijk (NL) (a) The sample and fibre alignment setup are on a table in front of the Co60 source. A laser, detector and measurement computer are outside of the radiation room. (b) Detail of the setup. The sample was mounted vertically to obtain a uniform absorbed dose over the sample surface. A dosimeter was put near the sample, and an Al plate is used as a buildup material.

dose over the lifetime. During irradiation, with each experiment taking a few days, the components were monitored by transmission measurements using a tunable laser and detector. The experimental setup is shown in Figure 9.13. On the optical side, the sample is set up on a table with on alignment stages, with a lensed fibres for input and output in a first experiment (Sept. 2004) and standard fibre with vertical fibre couplers in a second experiment (Feb. 2005). Fibres run through dog-legged holes in the wall between the radiation and control rooms. In the control room, a tunable laser, detector and computer are used to perform transmission measurements. The measurement setup is put in front of the radiation source and shielded from electrons[6] by an Al plate. A dosimeter is set up at the same distance from the source as the sample, also shielded by the Al plate and is calibrated to water. In order to convert the absorbed dose from krad (water) to krad (Si), a calibration factor of 0.896 must be applied.

The experimental setup had some problems. In the first experiment, facet coupling with lensed fibres was used, and the setup was put on a not very stable table that needed to be moved closer to the source after alignment. Also, the radiation room is located just near a (low traffic) road. This made alignment rather difficult and over the duration of the experiment (a few days), the alignment slowly degraded. In the second experiment, vertical fibre couplers were used which are more tolerant to misalignment. Still, the gradual degradation of alignment made it impossible to measure a possible radiation-induced insertion loss due to absorption. The Q factor of the resonators tested contains that information too, however. A second problem was the temperature control. In the first experiments, temperature control was not done. However, ambient temperature changes were relatively small. In the second series of experiments, a peltier element and thermistor with a controller in the control room were used, but the controller frequently failed for no obvious reason (possibilities include a short, electrical grounding of the setup or influence of radiation on the thermistor or peltier element). The temperature difference between day and night is clearly visible in the measurement results.

9.3.3 Results

A first measurement was performed on the drop port of a racetrack resonator with a 13 to 16 nm FSR. Figure 9.14 shows the resonance wavelength shift as a function of runtime and dose for five resonances. The dose rate was 82.1 rad (Si)/min. After 300 krad, the shift is about 100 pm.

---

[6]Each γ source, including Co60, also produces electrons.

Figure 9.14: Wavelength shift as a function of time and dose for 4 resonances of a racetrack resonator

However, the experiment was performed without temperature control. Room temperature (indicated on Fig. 9.14) did not change for more than 0.4°C, corresponding with an expected wavelength shift of 25 to 40 pm. Both the fluctuations in the optical signal with a typical period of 1 to 2 hours and the fluctuation with a period of 24 hours correspond to similar temperature fluctuations. Subtracting the total shift introduced by temperature drift from the measurement, this first measurement gives a maximum total radiation induced shift of 60 pm to 90 pm or 0.2 to 0.3 pm/krad. This corresponds with a core refractive index change of the order of 5 to $10 \cdot 10^{-7}$/krad

In a second experiment, a racetrack resonator with a 12 to 13 nm FSR was tested, but now with vertical fibre couplers for enhanced stability and with temperature control. The result is shown in Figure 9.15. A very linear wavelength shift of 100 to 120 pm was measured over a 296.576 krad (Si) dose (48 hour experiment), giving a shift of 0.314 ± 0.003 to 0.387 ± 0.002 pm/krad, which is of the same order as in the first experiment. The small difference with the first experiment can be explained by a difference in waveguide width and height. After 48 hours irradiation was stopped and the material relaxed partially and slowly.

During further experiments, the temperature controller failed, superimposing a fluctuation of tens to several tens of picometre onto the irradiation-induced wavelength shift. Even so, the roughly fitted shifts were of the same order as the previously measured ones.

Figure 9.15: Wavelength shift as a function of time and dose for 5 resonances of a racetrack resonator (temperature controlled)

No measurable degradation of the quality factors of the rings could be measured. Therefore, the increase of absorption accompanying the refractive index change seems to be very small.

Comparison with SoS device

In order to assess the radiation hardness of commercially available WDM devices, a SoS AWG demultiplexer device with a 100 GHz channel spacing was irradiated and monitored. The chip was fibre pigtailed and could easily be put close to the radiation source, limiting the duration of the experiment to 30 hours at a mean dose rate of 156.8 rad (Si)/min. Three channels of the AWG were monitored. Measurements on glasses have shown refractive index changes of the order of 2×10^{-8}/krad [103]. Given this index change, a wavelength shift of the order of 0.02 pm/krad can be expected in SoS waveguides, influenced in a non-trivial way however by the exact glass composition and the doping of the waveguide core. The measurements show a maximum shift of 0.03 pm/krad, which corresponds to a refractive index change of 2.8×10^{-8}/krad, which is of the same order as the previous measurements. With a 800pm channel spacing, this effect only becomes important at high TID and will give only a slight degradation of the bit error rate.

An important fact is that the measured shifts can easily be compensated for by temperature tuning. Therefore, even though the spectral shift of SOI components is an order of magnitude larger than that of glass based

components, it is no real problem as the devices will have to be temperature controlled anyway.

9.4 Conclusions

Already with the current crosstalk performance, real-world communication applications are possible for SOI nanophotonic devices, albeit with limited scalability to a few communicating nodes. One issue is the polarisation dependency, unless the application allows use of polarisation maintaining fibres. Polarisation diversity is a solution, but only if both circuits can be made exactly identical. Nevertheless, for specific applications, polarisation diversity may however now already deliver polarisation independent components. A polarisation-diversity AWG duplexer has been demonstrated, by looping both paths through the same AWG in opposite direction. This AWG was designed as a duplexer for a fibre-to-the-home application.

Another possible application is that of high-speed board-to-board interconnect networks. SOI components can deliver highly functional components, currently scalable to a limited number of communicating nodes. While use of SOI components for such a high-speed interconnect in space may not sound obvious, it actually very well meets the size, power and weight constraints of space. A demonstrator for such an application was built and tested.

An initial study of the radiation hardness has resulted γ radiation induced spectral shifts that can easily be compensated for by temperature tuning. Further radiation testing could lead to a better knowledge of the optical effects of radiation-induced changes in silicon-on-insulator. Neutron irradiation tests should be performed as these can be expected to have a different influence than γ tests, as neutron impact tends to create clusters of defects, which can enhance Rayleigh scattering and absorption. However, neutron irradiation is more cumbersome as it activates the samples.

An AWG component was fibre pigtailed using a fibre array connector and UV curable glue, which resulted in a mechanically attractive form of pigtailing. Vertical coupling, combined with the small die size, clearly opens up ways of packaging that are robust to vibrations, which is of primary importance for space and end user applications and eases handling. Further study of the mechanical properties of vertical coupling is needed however.

Chapter 10

Conclusions

10.1 Overview

High refractive index contrast, silicon-on-insulator (SOI) photonic wires allow for a dramatic size reduction of passive integrated wavelength filters for wavelength division multiplexing (WDM) applications. The main advantage of silicon is that the high-end tools and processes for fabrication of electronic integrated circuits can be used. In this work, characteristics of wavelength filters in silicon-on-insulator were studied. These filters were fabricated with CMOS tools and CMOS-based wafer-scale processes. These filters act as channel drop filters or wavelength routers. Filters for up to 16 channels were demonstrated with footprints as small as $0.01\,\text{mm}^2$ to $0.1\,\text{mm}^2$. This allows for many devices to be put in parallel on a single chip, or for multiple functions to be integrated in a single light-path on a single photonic integrated circuit.

Single ring resonator add-drop filters with bend radii of 1 to 5 µm have been fabricated, in a racetrack or bend-coupled configuration to enhance coupling. With this enhanced coupling, cavity losses are compensated for and the add-drop filters show nice extinction ratios with bandwidths compatible with high-speed (10Gbps) communication, while keeping the FSR as large as 10 nm or higher. Circular ring resonators showed Q factors higher than 10000. A second-order ring lattice filter was fabricated, but its transmission is currently limited by the spectral accuracy of the resonances. A third-order Butterworth filter based on a Mach-Zehnder and an all-pass ring resonator showed good extinction in one of the two output ports.

Non-resonant higher-order filters proved to be more successful with the current state of technology. Mach-Zehnder lattice channel drop filters with WDM channel spacings were designed and fabricated, showing

up to 15 dB isolation between output ports, with large passband widths. Up to an eleventh-order filter was demonstrated. The main limitations to crosstalk are phase errors in the delay lines, and coupler deviations, both in amplitude and phase. The waveguides forming the delay lines were made broader in order to reduce phase errors. In the bends and couplers, the waveguides are kept narrow to enforce a single propagating mode.

Several generations of AWG devices were designed and fabricated, and ultimately delivered channel isolation figures of 15 dB and more. The insertion loss was reduces through a double-etch process, locally lowering the refractive index contrast around the star couplers. Again, the waveguides were expanded to 800 nm in order to reduce phase errors. Nevertheless, phase errors still limit the crosstalk level, possibly together with unoptimized coupling between the arrayed waveguide apertures. Nevertheless, a 16-channel AWG shows a crosstalk lower than -15 dB and still fits within a 0.1 mm^2 footprint.

The phase errors, coupler variations and other fabrication variations limit the scalability of silicon-on-insulator components in a communication application. For applications with relatively larger power budgets, such as optical interconnects, the designed and fabricated filters may to scale to 16 channels with the current state of technology, with a modest WDM channel spacing 200 GHz, or 100 GHz for smaller channel numbers. For operation in telecommunication, for more power-constrained applications, or for a large number of channels or small channel spacings, sub-nanometre accuracy on the waveguide and structure dimensions is necessary. Several sources however currently point to a stochastic variation of the waveguide width the order of a couple of nanometres on a length scale of the order of half a micrometre: transmission measurements of filters combined with simulation, top-down microscope measurements, data available from fabrication reports of the CMOS reticles used for lithography and available literature.

On the other hand, waveguide and component loss as such is not a major issue anymore. Filters show a loss comparable to that of lower-index-contrast devices, when properly designed. Still, waveguide losses should be reduced further to enable higher channel isolations, and a higher tolerance to fabrication variations.

Apart from the stochastic width fluctuations, other sources of variations exist. The uniformity and accuracy of structures was characterized. The main influence of variations between nominally identical components is a shift of their transmission spectrum. We have seen this shift clearly depends on the distance between the components. For two identical struc-

tures next to each other, with a spacing of a few tens of micrometre, the shift is limited to about $\pm 5\,$nm or better. For structures spaced a few mm apart, the shift is larger, up to $\pm 10\,$nm. This can in part be explained by silicon layer thickness variations on that scale, and in part by larger width and sidewall slope variations between structures spaced over a longer distance. The die-to-die uniformity was measured in order to check for deterministic components relating to layer thickness or fabrication process dependencies on the position. The measured correlation between the position on the wafer and the relative spectral shifts are non-zero but small. No clear relation could therefore be seen. This points again to random fluctuations changing the mean waveguide width and dominating the spectral shift.

Despite the relatively high crosstalk, a silicon-on-insulator arrayed waveguide grating was used as a wavelength router in a demonstrator for a high-speed reconfigurable optical interconnect application. A board-to-board interconnect based on tuneable lasers, fixed receivers and passive wavelength routing, operated at 10 Gbps with a silicon-on-insulator AWG as the central component. This cyclic wavelength router had 4 inputs and 4 outputs, a -12 dB crosstalk level, and was fibre-pigtailed using vertical fibre coupling and a commercially available standard single-mode fibre array connector. Only at reduced input power, the device showed its scalability limitations. Careful extrapolation of the measured bit error rate for given optical powers showed a scalability of up to 9 ports. As the goal of the study was the assessment of optical interconnects for space applications, initial radiation hardness testing of silicon-on-insulator components was done. The devices tested showed a irradiation-induced spectral shift that is an order of magnitude larger than that of glass-based components, but can still be very well compensated for through temperature tuning.

Polarisation dependency is a major bottleneck for the use of high-index contrast devices. Polarisation diversity is probably the only viable solution. However, due to fabrication variations, this is not easily accomplished. Nevertheless, an AWG component without polarisation dependent spectral shift and with a polarisation dependent loss of only 0.7 dB was demonstrated. This was done by looping the two paths corresponding to the orthogonal fibre polarisations through the same AWG, but in reverse direction. The component was an initial study of a duplexer for fibre-to-the-home applications.

10.2 Conclusion

Wavelength division multiplexing filters were designed and fabricated in silicon-on-insulator with CMOS tools and processes, including 248 nm deep UV lithography. The devices have WDM channel spacings and a footprint of only 0.01 to 0.1 mm². Channel isolation levels are currently in the 12 to 20 dB range, limited by fabrication variations such as stochastic phase errors and deterministic and stochastic coupler deviations. Both for lattice filters and for arrayed waveguide grating devices, these are the best values reported to date for devices in this technology, with the unique added value of fabrication with CMOS tools and processes. In the context of a high-speed optical interconnect application, a fibre pigtailed AWG was demonstrated for the first time, and an initial study of the high-energy radiation hardness was performed.

Through testing in a demonstrator for this high-speed interconnect, the scalability of the studied components can be estimated at 4 to 8 ports, and for the best components (20 dB isolation) up to 16 channels, at least if the fibre coupling loss can be kept under control.

Analysis of the fabrication variations through measurement and simulation shows that a dimensional control better than a nanometre is needed, at least for variations on a length scale of the order of half a micrometre to several micrometre. Measurement and analysis of the non-uniformity of components at short and longer distances within a die and between dies, showed that the stochastic variations currently induce spectral shifts between nominally identical components of the order of ± 5 nm on a short length scale and twice that on a longer length scale.

Both the variations within a structure and the variations between structures can be compensated for, for instance through temperature tuning. However, with the current magnitude of non-uniformity, the required power would be high. Also, the multiple sources of inaccuracy would necessitate complex tuning algorithms if a multitude of components would need to be aligned. Furthermore, some problems such as backscattering and coupler deviations cannot be compensated for by tuning or are at least much harder to compensate for. Therefore, a highly parallel integration, with many identical components operating on parallel wavelength channels but with a limited functionality, seems a more viable kind of application in the near future.

In conclusion, this work has demonstrated ultra-compact SOI WDM filters, fabricated with CMOS based wafer-scale technology. At the current performance levels, it is possible to use SOI devices for communication

applications that only need a limited scalability and have a power budget of the order of 20 dB.

10.3   Recommendations

Engineering ultra-compact SOI wavelength filter devices further should tackle several points:

- **Improvement of the fabrication process** in order to reduce the crosstalk level, the accuracy and uniformity both on a short length scale and between dies, and the complexity of tuning mechanisms needed. Waveguide dimension variations should be reduced to the sub-nanometre level, at least on a length scale of several hundreds of micrometre and longer. Larger process windows are needed to be able to actually combine structures with different dimensions and with (quasi-periodic) structures of different pitches and fill factors. This is only possible to a limited extent with the current process.

- **CMOS fabrication oriented design of components**, where one tries to avoid the need for exact dimensional control, could enhance performance without the need for adding complex tuning.

- **Demonstration of polarisation independent devices**, either through extremely good fabrication technology, or more realistically, through polarisation diversity. This has not yet been demonstrated generically and is a major bottleneck for many applications. The requirement is again a better fabrication technology, combined with tuning. The approach used in chapter 9 to loop both polarisations in reverse direction through the same component, can work in cases with a limited functionality in a single light path.

- **Demonstration of robust fibre array coupling** is another major bottleneck. Highly efficient coupling to fibre has been demonstrated, and devices coupled to a very limited number of fibres (two) is setting its first commercial steps. However fibre coupling to an array of fibres, robust enough for real-world application, has not been shown. A first example of such fibre coupling was given in this work.

In other words, while applications with some degree of performance and scalability are possible today, possibly even on a commercial, high-volume scale through pervasive tuning of components, a lot of work is needed to make general silicon photonic integrated circuits reality.

Appendix A

Stochastic waveguide width deviations

A.1 Phase error model

A simple one-dimensional model can be used to simulate the effect on phase of small variations along the sidewall. The variations on each sidewall are considered to be stationary, and uncorrelated for both sidewalls. Along a sidewall, the points are considered to be highly correlated within a correlation length L_c. Within L_c, the sidewall approximately keeps a constant lateral position. The waveguide is therefore modelled as a series of equal-length chunks with differing sidewall positions. This is illustrated in figure A.1. The waveguide is supposed to be single mode and lossless, without coupling between TE and TM modes. As roughness-type variations give rise to scattering, the model is only valid for long enough L_c. Additionally, the stochastic process is considered to be stationary and ergodic. Often, we cannot expect the stochastic process to be stationary, let alone ergodic, over a longer length scale. However, here stationarity and ergodicity will be supposed, to model variations over relatively short length scales on a chip.

The phase difference after propagation through the corrugated waveguide, relative to the phase after propagation through the ideal waveguide with propagation constant β and length L, is given by

$$\delta\phi = \int_0^L \beta(x) \cdot dx - \beta \cdot L \cong \int_0^L \frac{\partial \beta}{\partial w} \cdot \delta w(x) \cdot dx, \qquad (A.1)$$

with $\beta(x)$ and $w(x)$ the local propagation constant and width, respectively.

Figure A.1: Phase error model

Using the discretized model, this is rewritten as

$$
\delta\phi = \begin{cases}
\frac{\partial\beta}{\partial w} \cdot \delta w \cdot L & , \quad L < L_c \\[2ex]
\sum_{i=1}^{L/L_c} \frac{\partial\beta}{\partial w} \cdot \delta w\,[i] \cdot L_c & , \quad L > L_c
\end{cases}
\tag{A.2}
$$

The variance of the phase error is then given by

$$
\sigma^2\left(\delta\phi\right) = \begin{cases}
\left(\frac{\partial\beta}{\partial w}\right)^2 \cdot \sigma^2\left(\delta w\,[i]\right) \cdot L^2 & , \quad L < L_c \\[2ex]
\left(\frac{\partial\beta}{\partial w}\right)^2 \cdot L_c^2 \cdot \sum_{i=1}^{L/L_c} \sigma^2\left(\delta w\,[i]\right) & , \quad L > L_c
\end{cases}
\tag{A.3}
$$

In the second equation of A.3, the variance of the sum was replaced by the sum of the variances in each section, as on a distance longer than the correlation length, the points are considered uncorrelated. Because the distribution of $\delta w\,[i]$ is the same in each section, $\sigma^2\left(\delta w\,[i]\right) \equiv \sigma^2\left(\delta w\right)$. While the distributions of the variation of both sidewalls are the same, the variations are considered to be uncorrelated. Therefore, the variance of the waveguide width deviation is given by $\sigma_{\delta w}^2 = 2 \cdot \sigma_{d(sidewall)}^2$. The variance of the phase error finally reduces to

$$
\sigma^2\left(\delta\phi\right) = \begin{cases}
2 \cdot \left(\frac{\partial\beta}{\partial w}\right)^2 \cdot \sigma_{d(sidewall)}^2 \cdot L^2 & , \quad L < L_c \\[2ex]
2 \cdot \left(\frac{\partial\beta}{\partial w}\right)^2 \cdot \sigma_{d(sidewall)}^2 \cdot L_c \cdot L & , \quad L > L_c
\end{cases}
\tag{A.4}
$$

With the stochastic process ergodic, the ensemble variance of the phase error equals the variance taken along the wire.

A.2   Mean width

What is the relation between the mean wire width and the parameters describing the stochastic variations of the wire sidewall ? Suppose the stochastic process describing the sidewalls is stationary and ergodic. Then, for an infinitely long wire, the expected value of the mean sidewall position is equal to the mean at each point, the nominal sidewall position. The variance over an ensemble of wires goes to zero. However, in filters, the wires have a finite, and often even quite short length.

Due to the finite length, the (spatial) mean width and the spatial variance of mean width will be different from zero. Due to ergodicity, the ensemble mean and variance will be equal to the spatial mean and variance. We take the simplified model from A.1 again as an approximation. The width in a section i is given by $w_i = w + d(sidewall1)_i + d(sidewall2)_i$, with w the nominal (designed) width.

The spatial mean along the wire is then

$$
\begin{aligned}
\mu_w^x &= \frac{1}{L/L_c} \sum_{i=1}^{L/L_c} w(x) \\
&= \frac{L_c}{L} \sum_{i=1}^{L/L_c} \left(w + d(sidewall1)_i + d(sidewall2)_i \right) \\
&= w + \frac{L_c}{L} \left(\sum_{i=1}^{L/L_c} d(sidewall1)_i + \sum_{i=1}^{L/L_c} d(sidewall2)_i \right)
\end{aligned}
\tag{A.5}
$$

So, for $L \to \infty$, then $\mu_w^x \to w$. Due to ergodicity then $E[\mu_w] = \mu_w^x$.

It can be confusing to think that because the expected value of the mean width approaches the nominal width for infinite length, the phase error would go to zero. The phase error could indeed be written in terms of the mean width, but the propagation is still over the total length L, which is not taken into account when calculating the mean width. So, the limit for $L \to \infty$ of

$$
\begin{aligned}
\delta\phi &= (w - \mu_w^x) \frac{d\beta}{dw} L \\
. &= \frac{d\beta}{dw} \left[wL - L_c \left(\sum_{i=1}^{L/L_c} d(sidewall1)_i + \sum_{i=1}^{L/L_c} d(sidewall2)_i \right) \right]
\end{aligned}
\tag{A.6}
$$

is not straightforward. Calculating the spatial variance (equal to the ensemble variance) of eq. A.6 leads again to eq. A.4.

Taking the variance of eq. A.5 gives

$$\sigma_{\mu_w^x}^2 = \frac{L_c^2}{L^2} \left[\sigma^2 \left(\sum_{i=1}^{L/L_c} d(sidewall1)_i \right) + \sigma^2 \left(\sum_{i=1}^{L/L_c} d(sidewall2)_i \right) \right]$$

$$= 2\frac{L_c}{L}\sigma_{d(sidewall)}^2, \tag{A.7}$$

where we used again the fact that at distances longer that the sections along the sidewall are uncorrelated and both sidewalls are uncorrelated. Figure A.2 shows the standard deviation of the mean width as a function of the wire length and the σ of the one-dimensional sidewall position variations, for three different correlation lengths. Note the model is actually not valid for L not sufficiently larger than L_c.

(a) L$_c$=100 nm

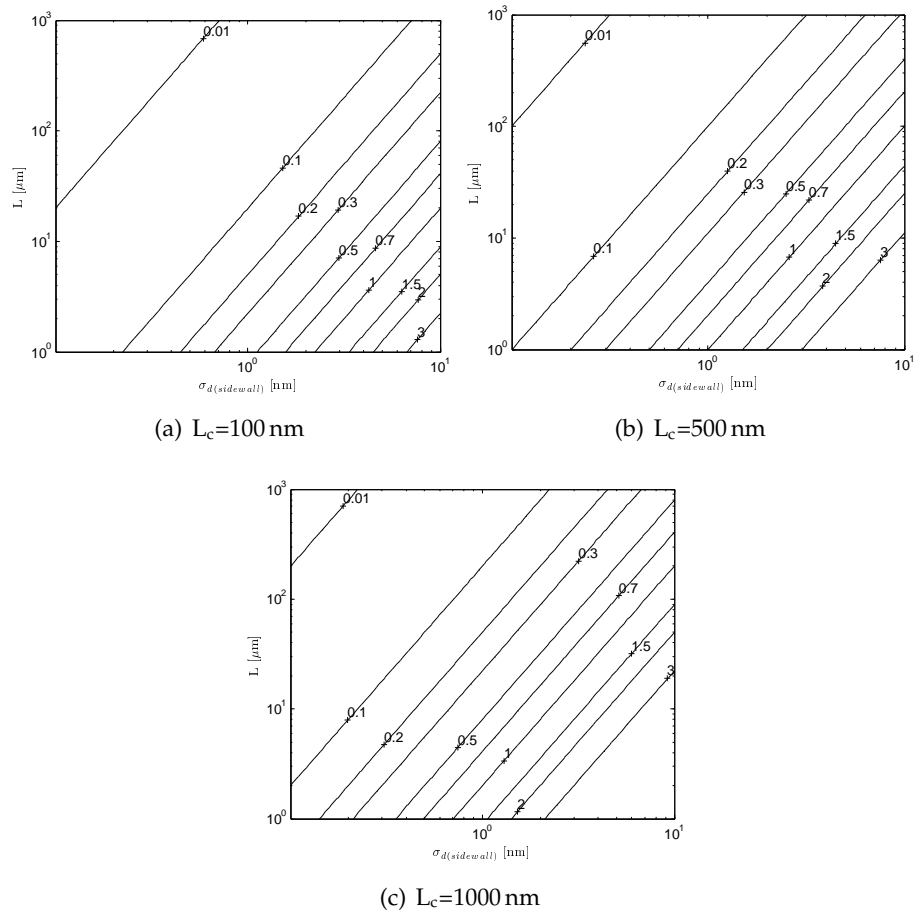(b) L$_c$=500 nm

(c) L$_c$=1000 nm

Figure A.2: Calculated equicontours of the standard deviation of the mean width in nm as a function of the standard deviation of the sidewall position and the length of the wire, for different correlation lengths.

Appendix B

Designed masks and fabricated components

During the course of this work, many masks with test structures were designed. The full reticles gathered masks designed by multiple partners within the research group and from other European and international research groups. Here, an overview of the contributions to the masks is given, with a brief listing of the components relevant to this work. During 2002-2005, the masks were named after the international projects in which contexts the masks were at least partly designed: the FP4 PICCO project and the FP5 PICMOS project. During the PICCO project, four reticles were designed (PICCO_XX), three of which this work contributed to. The fifth and sixth CMOS reticles were partly supported by the PICMOS project (PICMOS_XX). From the 7th reticle on, the name PICSOI_XX was adopted Each reticle is divided in cells, with various names such as A1,A2,.... These names are indicated below. Multiple mask structures, especially the AWG devices, were designed by Dr. Wim Bogaerts.

PICCO_02 (Q1 2002)

- simple wire structures:long linear tapers, wires of different lengths, directional couplers, single s-bends, mirror bends, straight-to-bend directional couplers

- ring resonators: ring resonators, parallel ring resonators, 2nd order series ring resonators, polygonal rings

- photonic crystal waveguides: W1, trench waveguides, 1 changed hole,

PICCO_03 (Q1 2003)

This mask contained a cell with photonic crystal waveguides: W1,W3,one hole shifted,W1 with longer border holes,hybrid wire/PhC, s-bends, hybrid bends, splitters

PICCO_04 (Q3 2003)

- A6

  - wire structures: wires (up to 2mm length), angled crossings, y junctions, horizontally and vertically folded waveguides, spirals (up to 10mm length)
  - ring resonators: different radiuses, ring and racetrack, symmetric and asymmetric coupling, bend coupling
  - simple ring resonator demux
  - rectangular resonators, double rectangular resonators
  - mach-zehnders with bends in 1 arm, varying delay length and coupling length (interleaving)
  - lattice filters with 5 stages, discrete set of coupling lengths, cascades of lattice filters

- A7 First AWG designs, with classical layout (large bend radius)

PICMOS_01 (Q4 2004)

- A11: spirals (up to 25mm length), horizontally folded waveguides, 1x1 symmetrical and asymmetrical MZI, tapered crossings

- A12: OPC test structures, Y junctions, 1x1 and 1x2 MMIs (tapered), directional couplers (different radiuses and lengths, bend-to-bend and straight-to-bend), all-pass ring resonators,

- A14: y junctions, splitter trees, 2nd generation AWG designs

- ROUTER: ESA multigigabit optical backplane project components

PICMOS_02 (Q2 2005)

- A16: ring resonators, all-pass ring resonators, add-drop ring resonators, MZI with single ring all-pass filters in one arm, Mach-Zehnder (MZI) with delay length, MZI as coupler, Mach-Zehnder cascade

(channel selectors), Mach-Zehnder lattice filters, ring lattice filters, simple ring demux

- ROUTER2: first shallowly etched AWG designs, both in ESA project format and as single components

PICSOI_07 (Q4 2005)

- ROUTER3: shallow and deep-shallow AWG designs (both in ESA project format and as single components), polarisation diversity circuits, shallow crossings, AWG with heaters, 2x2 MZI as switches (with heaters), 2x2 crossconnect

PICSOI_08 (Q2 2006)

- ROUTER4: AWG designs (shallow, deep-shallow, grating aperture angle variation,...), AWG-based polarisation independent duplexer.

- EPIX3: 2x2 MZI, 1x2 MZI with all-pass filter, single ring add-drop filters, series of identical all-pass filters on single wire, directional couplers with large bend radius

Appendix C

Publications

C.1 International Journals

Publications directly originating from the work described in this thesis:

1. P. Bienstman, L. Vanholme, W. Bogaerts, P. Dumon, P. Vandersteegen, "Python in Nanophotonics Research", accepted for publication in *Computing in Science and Engineering*, (to be published)

2. W. Bogaerts, D. Taillaert, P. Dumon, D. Van Thourhout, R. Baets, "A polarization-diversity wavelength duplexer circuit in silicon-on-insulator photonic wires",*Optics Express*, 15(4), p. 1567-1578 (2007)

3. W. Bogaerts, P. Dumon, D. Van Thourhout, D. Taillaert, P. Jaenen, J. Wouters, S. Beckx, R. Baets, "Compact Wavelength-Selective Functions in Silicon-on-Insulator Photonic Wires",*J. Selected Topics in Quantum Electronics*, 12(6), p. 1394-1401 (2006)

4. P. Dumon, G.R.A. Priem, L.R. Nunes, W. Bogaerts, D. Van Thourhout, P. Bienstman, T.K. Liang, M. Tsuchiya, P. Jaenen, S. Beckx, J. Wouters, R. Baets, "Linear and nonlinear nanophotonic devices based on silicon-on-insulator wire waveguides", *Japanese Journal of Applied Physics*, 45(8B), p.6589-6602(2006)

5. P. Dumon, W. Bogaerts, D. Van Thourhout, D. Taillaert, R. Baets, J. Wouters, S. Beckx, P. Jaenen, "Compact wavelength router based on a Silicon-on-insulator arrayed waveguide grating pigtailed to a fiber array", *Optics Express*, 14(2), p.664-669 (2006)

6. W. Bogaerts, R. Baets, P. Dumon, V. Wiaux, S. Beckx, D. Taillaert, B. Luyssaert, J. Van Campenhout, P. Bienstman, D. Van Thourhout,

Nanophotonic Waveguides in Silicon-on-Insulator Fabricated with CMOS Technology, *Journal of Lightwave Technology* (invited), 23(1), p.401-412 (2005)

7. P. Dumon, W. Bogaerts, V. Wiaux, J. Wouters, S. Beckx, J. Van Campenhout, D. Taillaert, B. Luyssaert, P. Bienstman, D. Van Thourhout, R. Baets, Low-loss SOI Photonic Wires and Ring Resonators Fabricated with Deep UV Lithography,*Photonics Technology Letters*, 16(5), p.1328-1330 (2004)

8. W. Bogaerts, D. Taillaert, B. Luyssaert, P. Dumon, J. Van Campenhout, P. Bienstman, D. Van Thourhout, R. Baets, V. Wiaux, S. Beckx, Basic structures for photonic integrated circuits in Silicon-on-insulator,*Optics Express* (invited), 12(8), p.1583-1591 (2004)

9. W. Bogaerts, P. Dumon, D. Taillaert, V. Wiaux, S. Beckx, B. Luyssaert, J. Van Campenhout, D. Van Thourhout, R. Baets, SOI Nanophotonic Waveguide Structures Fabricated with Deep UV Lithography , *Photonics and Nanostructures: Fundamentals and Applications* (invited), 2(2), p.81-86 (2004)

Publications in the context of collaboration and the facility access service:

1. T.K. Liang, L.R. Nunes, M. Tuschiya, K.S. Abedin, T. Miyazaki, D. Van Thourhout, W. Bogaerts, P. Dumon, R. Baets, H.K. Tsang, "High speed logic gate using two-photon absorption in Silicon waveguides", *Optics Communications*, 265(1), p. 171-174 (2006)

2. G. Roelkens, J. Brouckaert, D. Taillaert, P. Dumon, W. Bogaerts, D. Van Thourhout, R. Baets, Integration of InP/InGaAsP photodetectors onto Silicon-on-Insulator waveguide circuits, *Optics Express*, 13(25), p.10102-10108 (2005)

3. G. Roelkens, P. Dumon, W. Bogaerts, D. Van Thourhout, R. Baets, Efficient Silicon-on-Insulator fiber coupler fabricated using 248nm deep UV lithography, *Photonics Technology Letters*, 17(12), p.2613-2615 (2005)

4. G.R.A. Priem, P. Dumon, W. Bogaerts, D. Van Thourhout, G. Morthier, R. Baets, Optical bistability and pulsating behaviour in Silicon-On-Insulator ring resonator structures, *Optics Express*, 13(23), p.9623 - 9628 (2005)

5. T. K. Liang, L. R. Nunes, T. Sakamoto, K. Sasagawa, T. Kawanishi, M. Tsuchiya, G.R.A. Priem, D. Van Thourhout, P. Dumon, R. Baets, H. K. Tsang, Ultrafast all-optical switching by cross-absorption modulation in silicon wire waveguides, *Optics Express*, 13(19), p.7298 (2005)

6. B. Luyssaert, P. Bienstman, P. Vandersteegen, P. Dumon, R. Baets, Efficient Nonadiabatic Planar Waveguide Tapers, *Journal of Lightwave Technology*, 23(8), p.2462-2468 (2005)

7. P. Sanchis, J. Marti, B. Luyssaert, P. Dumon, P. Bienstman, R. Baets, Analysis and design of efficient coupling in photonic crystal circuits, *Optical and Quantum Electronics*, 37(1-3), p.133-147 (2005)

8. P. Sanchis, J. Marti, W. Bogaerts, P. Dumon, D. Van Thourhout, R. Baets, Experimental results on adiabatic coupling into SOI photonic Crystal coupled-cavity waveguides, *Photonics Technology Letters*, 17(6), p.1199-1201 (2005)

9. B. Luyssaert, P. Vandersteegen, D. Taillaert, P. Dumon, W. Bogaerts, P. Bienstman, D. Van Thourhout, V. Wiaux, S. Beckx, R. Baets, A Compact Horizontal Spot-Size Converter Realized in Silicon-on-Insulator,*Photonics Technology Letters*, 17(1), p.73-75 (2005)

10. P.I. Borel, L.H. Frandsen, A. Harpoeth, J.B. Leon, H. Liu, M. Kristensen, W. Bogaerts, P. Dumon, R. Baets, V. Wiaux, S. Beckx, J. Wouters, Bandwidth engineering of photonic crystal waveguide bends, *Electronics Letters*, 40(20), p.1263-1264 (2004)

11. P. Sanchis, J. Garcia, J. Marti, W. Bogaerts, P. Dumon, D. Taillaert, R. Baets, V. Wiaux, J. Wouters, S. Beckx, Experimental demonstration of high coupling efficiency between wide ridge dielectric waveguides and single-mode photonic crystal waveguides, *Photonics Technology Letters*, 16(10), p.2272-2274 (2004)

12. L.H. Frandsen, P.I. Borel, Y.X. Zhuang, A. Harpoeth, M. Thorhauge, M. Kristensen, W. Bogaerts, P. Dumon, R. Baets, V. Wiaux, S. Beckx, J. Wouters, Ultra Low Loss 3 dB Photonic Crystal Waveguide Splitter ,*Optics Letters*, 29(14), p.1623-1625 (2004)

C.2   Book chapters

1. D. Van Thourhout, W. Bogaerts, P. Dumon, Sub-micron Silicon strip waveguidess, *Optical Interconnects: The Silicon Approach*, eds. L. Pavesi and G. Guillot, Springer (2006)

C.3   Selected Papers in International Conference Proceedings

These conference papers originate directly from this work.

1. E.D. Kyriakis-Bitzaros, E. Grivas, G. Halkias, S. Katsafouros, P. Dumon, G. Morthier, R. Baets, T. Farell, N. Ryan, I. McKenzie, E. Armadillo, A WDM Optical Backplane with AWG Based Passive Routing, *Photonics in Switching 2006*, Greece, p.273-275 (2006)

2. P. Bienstman, P. Dumon, W. Bogaerts, D. Taillaert, F. Van Laere, K. De Vos, D. Van Thourhout, R. Baets, Silicon nanophotonics using deep-UV lithography, *APOC 2006 (invited)*, 6351(99), South Korea, (2006)

3. D. Van Thourhout, W. Bogaerts, P. Dumon, G. Roelkens, J. Van Campenhout, R. Baets, Functional Silicon Wire Waveguides, *Integrated Photonics Research and Application, Nanophotonics, OSA Topical Meetings*, United States, p.IWA6.pfd (2006)

4. G.R.A. Priem, P. Dumon, W. Bogaerts, D. Van Thourhout, G. Morthier, R. Baets, Nonlinear effects in ultrasmall silicon-on-insulator ring resonators, *Photonics Europe*, France, p.6183-47 (2006)

5. P. Dumon, W. Bogaerts, D. Van Thourhout, G. Morthier, R. Baets, P. Jaenen, S. Beckx, J. Wouters, T. Farrell, N. Ryan, E. Grivas, E. Kyriakis-Bitzaros, G. Halkias, I. McKenzie, A nanophotonic 4x4 wavelength router in Silicon-on-insulator, *Optical Fiber Communication Conference (OFC) 2006*, United States, p.Oth05 (2006)

6. W. Bogaerts, P. Dumon, J. Wouters, P. Jaenen, S. Beckx, V. Wiaux, D. Van Thourhout, D. Taillaert, B. Luyssaert, R. Baets, Silicon-on-Insulator Nanophotonics, *SPIE Optics and Optoelectronics (invited)*, 5956, Poland, p.5956R-1-15 (2005)

7. G.R.A. Priem, P. Dumon, W. Bogaerts, D. Van Thourhout, G. Morthier, R. Baets, Optical bistability analysis inside a two-bus ring resonator, *LEOS Annual Meeting 2005*, Australia, p.MN4 (2005)

8. D. Van Thourhout, P. Dumon, W. Bogaerts, G. Roelkens, G.R.A. Priem, R. Baets, Recent Progress in SOI Nanophotonic Waveguides, *European Conference on Optical Communication (ECOC) 2005 (invited)*, United Kingdom, p.Tu3.6.2 (2005)

9. P. Dumon, G. Roelkens, W. Bogaerts, D. Van Thourhout, J. Wouters, S. Beckx, P. Jaenen, R. Baets, Basic Photonic Wire Components in Silicon-on-Insulator, *Group IV Photonics*, Belgium, p.189-191 (2005)

10. P. Dumon, R. Baets, R. Kappeler, D. Barros, I. McKenzie, D. Doyle, Measured radiation sensitivity of Silica-on-Silicon and Silicon-on-insulator micro-photonic devices for potential space application,Optics and Photonics, *Photonics for Space Environments X*, 5897, United States, (2005)

11. P. Dumon, I. Christiaens, W. Bogaerts, V. Wiaux, J. Wouters, S. Beckx, D. Van Thourhout, R. Baets, Microring resonators in Silicon-on-insulator, *ECIO 2005*, France, p.196-199 (2005)

12. P. Dumon, W. Bogaerts, D. Van Thourhout, D. Taillaert, V. Wiaux, S. Beckx, J. Wouters, R. Baets, Wavelength-selective components in SOI photonic wires fabricated with deep UV lithography, *Group IV Photonics*, Hong Kong, p.WB5 (2004)

13. P. Dumon, W. Bogaerts, D. Van Thourhout, D. Taillaert, V. Wiaux, S. Beckx, J. Wouters, R. Baets, Cascaded Mach-Zehnder Filters in Silicon-on-Insulator Photonic Wires fabricated with deep UV lithography, *European Conference on Optical Communication (ECOC) 2004*, Th.(Post-deadline), Sweden, p.24-25 (2004)

14. P. Dumon, W. Bogaerts, J. Van Campenhout, V. Wiaux, J. Wouters, S. Beckx, R. Baets, Low-loss photonic wires and compact ring resonators in Silicon-on-insulator, *SPIE International Symposium Photonics Europe*, France, (2004)

15. P. Dumon, W. Bogaerts, J. Van Campenhout, V. Wiaux, J. Wouters, S. Beckx, R. Baets, Low-loss, single-mode photonic wires ring resonators in Silicon-on-Insulator, *LEOS 2003*, United States, p.289-299 (2003)

Bibliography

[1] W. Bogaerts. *Nanophotonic Waveguides and Photonic Crystals in Silicon-on-Insulator*. PhD thesis, Ghent University, 2004.

[2] W. Bogaerts, R. Baets, P. Dumon, V. Wiaux, S. Beckx, D. Taillaert, B. Luyssaert, J. Van Campenhout, P. Bienstman, and D. Van Thourhout. Nanophotonic waveguides in silicon-on-insulator fabricated with cmos technology. *J Lightwave Technol*, 23(1):401–412, Jan 2005.

[3] V. R. Almeida, Q. F. Xu, and M. Lipson. Ultrafast integrated semiconductor optical modulator based on the plasma-dispersion effect. *Opt Lett*, 30(18):2403–2405, Sep 15 2005.

[4] M. Rouviere, L. Vivien, X. Le Roux, J. Mangeney, P. Crozat, C. Hoarau, E. Cassan, D. Pascal, S. Laval, J. M. Fedeli, J. F. Damlencourt, J. M. Hartmann, and S. Kolev. Ultrahigh speed germanium-on-silicon-on-insulator photodetectors for 1.31 and 1.55 mu m operation. *Appl Phys Lett*, 87(23):–, Dec 5 2005.

[5] T. Tsuchizawa, K. Yamada, H. Fukuda, T. Watanabe, J. Takahashi, M. Takahashi, T. Shoji, E. Tamechika, S. Itabashi, and H. Morita. Microphotonics devices based on silicon microfabrication technology. *IEEE J. Sel .Topics Quant. Elec.*, 11(1):232–240, Jan-Feb 2005.

[6] P. Debackere, S. Scheerlinck, P. Bienstman, and R. Baets. Surface plasmon interferometer in silicon-on-insulator: novel concept for an integrated biosensor. *Opt Express*, 14(16):7063–7072, Aug 7 2006.

[7] R. Nagarajan, C. H. Joyner, R. P. Schneider, J. S. Bostak, T. Butrie, A. G. Dentai, V. G. Dominic, P. W. Evans, M. Kato, M. Kauffman, D. J. H. Lambert, S. K. Mathis, A. Mathur, R. H. Miles, M. L. Mitchell, M. J. Missey, S. Murthy, A. C. Nilsson, F. H. Peters, S. C. Pennypacker, J. L. Pleumeekers, R. A. Salvatore, R. K. Schlenker, R. B. Taylor, H. S. Tsai, M. F. Van Leeuwen, J. Webjorn, M. Ziari, D. Perkins,

J. Singh, S. G. Grubb, M. S. Reffle, D. G. Mehuys, F. A. Kish, and D. F. Welch. Large-scale photonic integrated circuits. *IEEE J. Sel .Topics Quant. Elec.*, 11(1):50–65, Jan-Feb 2005.

[8] E. D. Palik. *Handbook of Optical Constants of Solids*. Academic Press, 1985.

[9] J. A. McCaulley, V. M. Donnelly, M. Vernon, and I. Taha. Temperature-dependence of the near-infrared refractive-index of silicon, gallium-arsenide, and indium-phosphide. *Phys Rev B*, 49(11):7408–7417, Mar 15 1994.

[10] S. Olivier, C. Weisbuch, and H. Benisty. Compact and fault-tolerant photonic crystal add-drop filter. *Opt Lett*, 28(22):2246–2248, Nov 15 2003.

[11] T. Matsumoto, K. S. Eom, and T. Baba. Focusing of light by negative refraction in a photonic crystal slab superlens on silicon-on-insulator substrate. *Opt Lett*, 31(18):2786–2788, Sep 15 2006.

[12] L. Frandsen, P. Borel, J. S. Jensen, and O. Sigmund. Topology optimized photonic wire splitters. In *CLEO/QELS 2006*, page CMV4, Long Beach, CA, 2006.

[13] B. Luyssaert, P. Bienstman, P. Vandersteegen, P. Dumon, and R. Baets. Efficient nonadiabatic planar waveguide tapers. *J Lightwave Technol*, 23(8):2462–2468, Aug 2005.

[14] Fimmwave. Photon Design, http://www.photond.com.

[15] P. Bienstman. Two-stage mode finder for waveguides with a 2d cross-section. *Opt Quant Electron*, 36(1 - 3):5–14, 2004.

[16] D. Taillaert, R. Baets, P. Dumon, W. Bogaerts, D. V. Thourhout, B. Luyssaert, V. Wiaux, S. Beckx, and J. Wouters. Silicon-on-insulator platform for integrated wavelength-selective components. In *Workshop on Fibres and Optical Passive Components*, Palermo, Italy, June 2005.

[17] D. X. Dai and S. L. He. Analysis of multimode effects in the free-propagation region of a silicon-on-insulator-based arrayed-waveguide grating demultiplexer. *Appl Optics*, 42(24):4860–4866, Aug 20 2003.

[18] A. C. Turner, C. Manolatou, B. S. Schmidt, M. Lipson, M. A. Foster, J. E. Sharping, and A. L. Gaeta. Tailored anomalous group-velocity dispersion in silicon channel waveguides. *Opt Express*, 14(10):4357–4362, May 15 2006.

[19] E. Dulkeith, F. N. Xia, L. Schares, W. M. J. Green, and Y. A. Vlasov. Group index and group velocity dispersion in silicon-on-insulator photonic wires. *Opt Express*, 14(9):3853–3863, May 1 2006.

[20] G. P. Agrawal. *Nonlinear fiber optics*. Academic Press, San Diego, 2nd edition, 1995. Govind P. Agrawal. ill. ; 24 cm. Optics and photonics.

[21] W. R. Thurber, R. L. Mattis, Y. M. Liu, and J. J. Filliben. Resistivity-dopant density relationship for boron-doped silicon. *J Electrochem Soc*, 127(10):2291–2294, 1980.

[22] R. A. Soref and B. R. Bennett. Electrooptical effects in silicon. *IEEE J. Quantum Elect.*, 23(1):123–129, Jan 1987.

[23] M. Borselli, T. J. Johnson, and O. Painter. Measuring the role of surface chemistry in silicon microphotonics. *Appl Phys Lett*, 88(13):–, Mar 27 2006.

[24] G. Priem. *Nonlinear Behaviour in Nanophotonic Waveguides and Resonators for Ultrafast Signal Processing*. PhD thesis, Ghent University, 2006.

[25] W. Bogaerts, V. Wiaux, D. Taillaert, S. Beckx, B. Luyssaert, P. Bienstman, and R. Baets. Fabrication of photonic crystals in silicon-on-insulator using 248-nm deep uv lithography. *IEEE J. Sel .Topics Quant. Elec.*, 8(4):928–934, Jul-Aug 2002.

[26] T. Barwicz and H. A. Haus. Three-dimensional analysis of scattering losses due to sidewall roughness, in microphotonic waveguides. *J Lightwave Technol*, 23(9):2719–2732, Sep 2005.

[27] F. P. Payne and J. P. R. Lacey. A theoretical-analysis of scattering loss from planar optical wave-guides. *Opt Quant Electron*, 26(10):977–986, Oct 1994.

[28] Y. A. Vlasov and S. J. McNab. Losses in single-mode silicon-on-insulator strip waveguides and bends. *Opt Express*, 12(8):1622–1631, Apr 19 2004.

[29] V. R. Almeida, R. R. Panepucci, and M. Lipson. Nanotaper for compact mode conversion. *Opt Lett*, 28(15):1302–1304, Aug 1 2003.

[30] H. Fukuda, K. Yamada, T. Shoji, M. Takahashi, T. Tsuchizawa, T. Watanabe, J. Takahashi, and S. Itabashi. Four-wave mixing in silicon wire waveguides. *Opt Express*, 13(12):4629–4637, Jun 13 2005.

[31] A. Sakai, G. Hara, and T. Baba. Propagation characteristics of ultrahigh-delta optical waveguide on silicon-on-insulator substrate. *Jpn J Appl Phys 2*, 40(4B):L383–L385, Apr 15 2001.

[32] A. Sakai, T. Fukazawa, and T. Baba. Low loss ultra-small branches in a silicon photonic wire waveguide. *Ieice T Electron*, E85C(4):1033–1038, Apr 2002.

[33] D. K. Sparacin, S. J. Spector, and L. C. Kimerling. Silicon waveguide sidewall smoothing by wet chemical oxidation. *J Lightwave Technol*, 23(8):2455–2461, Aug 2005.

[34] R. Orobtchouk, N. Schell, T. Benyattou, and J. M. Fedeli. Compact building blocks for optical link on soi technology. In *ECIO*, Grenoble, France, April 2005 2005.

[35] R. L. Espinola, M. C. Tsai, J. T. Yardley, and R. M. Osgood. Fast and low-power thermooptic switch on thin silicon-on-insulator. *IEEE Photonic Techn. Lett.*, 15(10):1366–1368, Oct 2003.

[36] H. Yamada, T. Chu, S. Ishida, and Y. Arakawa. Optical directional coupler based on si-wire waveguides. *IEEE Photonic Techn. Lett.*, 17(3):585–587, Mar 2005.

[37] J. Niehusmann, A. Vorckel, P. H. Bolivar, T. Wahbink, W. Henschel, and H. Kurz. Ultrahigh-quality-factor silicon-on-insulator microring resonator. *Opt Lett*, 29(24):2861–2863, Dec 15 2004.

[38] W. N. Ye, D. X. Xu, S. Janz, P. Cheben, M. J. Picard, B. Lamontagne, and N. G. Tarr. Birefringence control using stress engineering in silicon-on-insulator (soi) waveguides. *J Lightwave Technol*, 23(3):1308–1318, Mar 2005.

[39] D. Taillaert, H. Chong, P. I. Borel, L. H. Frandsen, R. M. De La Rue, and R. Baets. A compact two-dimensional grating coupler used as a polarization splitter. *IEEE Photonic Techn. Lett.*, 15(9):1249–1251, Sep 2003.

[40] Biegelse.Dk. Photoelastic tensor of silicon and volume dependence of average gap. *Phys Rev Lett*, 32(21):1196–1199, 1974.

[41] V. Constantoudis, G. P. Patsis, A. Tserepi, and E. Gogolides. Quantification of line-edge roughness of photoresists. ii. scaling and fractal analysis and the best roughness descriptors. *J Vac Sci Technol B*, 21(3):1019–1026, May-Jun 2003.

[42] A. Yamaguchi and O. Komuro. Characterization of line edge roughness in resist patterns by using fourier analysis and auto-correlation function. *Japanese Journal of Applied Physics Part 1-Regular Papers Short Notes and Review Papers*, 42(6B):3763–3770, Jun 2003.

[43] T. Goh, S. Suzuki, and A. Sugita. Estimation of waveguide phase error in silica-based waveguides. *J Lightwave Technol*, 15(11):2107–2113, Nov 1997.

[44] L. Prkna, M. Hubalek, and J. Ctyroky. Field modeling of circular microresonators by film mode matching. *IEEE J. Sel .Topics Quant. Elec.*, 11(1):217–223, Jan-Feb 2005.

[45] A. Sakai, T. Fukazawa, and T. Baba. Estimation of polarization crosstalk at a micro-bend in si-photonic wire waveguide. *J Lightwave Technol*, 22(2):520–525, Feb 2004.

[46] T. Fukazawa, T. Hirano, F. Ohno, and T. Baba. Low loss intersection of si photonic wire waveguides. *Jpn J Appl Phys 1*, 43(2):646–647, Feb 2004.

[47] S. G. Johnson, C. Manolatou, S. H. Fan, P. R. Villeneuve, J. D. Joannopoulos, and H. A. Haus. Elimination of cross talk in waveguide intersections. *Opt Lett*, 23(23):1855–1857, Dec 1 1998.

[48] M. Settle, M. Salib, A. Michaeli, and T. F. Krauss. Low loss silicon on insulator photonic crystal waveguides made by 193nm optical lithography. *Opt Express*, 14(6):2440–2445, Mar 20 2006.

[49] C. X. Wang, X. S. Xu, F. Li, W. Du, G. G. Xiong, Y. L. Liu, and H. D. Chen. Light-propagation characteristics of photonic crystal waveguide based on soi materials at different polarized states. *Chinese Physics Letters*, 23(9):2472–2475, Sep 2006.

[50] H. J. Levinson. *Principles of lithography*. SPIE, Bellingham, 2001.

[51] H. Ito. Chemical amplification resists: History and development within ibm. *IBM Journal of Research and Development*, 41(1/2):69, 1997.

[52] T. Heins, U. Dersch, R. Liebe, and J. Richter. Line edge roughness on photo lithographic masks. In *Microlithography*, volume 6152, pages 6152–62, San Jose, CA, 2006.

[53] L. H. A. Leunissen, W. G. Lawrence, and M. Ercken. Line edge roughness: experimental results related to a two-parameter model. *Microelectron Eng*, 73-74:265–270, Jun 2004.

[54] T. Yamaguchi, K. Yamazaki, M. Nagase, and H. Namatsu. Line-edge roughness: Characterization and material origin. *Japanese Journal of Applied Physics Part 1-Regular Papers Short Notes and Review Papers*, 42(6B):3755–3762, Jun 2003.

[55] M. Williamson and A. Neureuther. Enhanced, quantitative analysis of resist image contrast upon line edge roughness (ler). In *SPIE Microlithography*, volume 5039, pages 423–432, Santa Clara, CA, 2003.

[56] H. Ito. Dissolution behavior of chemically amplified resist polymers for 248-,193-, and 157-nm lithography. *IBM Journal of Research and Development*, 45(5):683, 2001.

[57] D. Taillaert. *Grating couplers as Interface between Optical Fibers and Nanophotonic Waveguides*. PhD thesis, Ghent University, 2004.

[58] D. Munteanu, C. Maleville, S. Cristoloveanu, H. Moriceau, B. Aspar, C. Raynaud, O. Faynot, J. L. Pelloie, and A. J. AubertonHerve. Detailed characterization of unibond material. *Microelectron Eng*, 36(1-4):395–398, Jun 1997.

[59] M. Tajima, Z. Q. Li, and R. Shimidzu. Photoluminescence mapping system applicable to 300 mm silicon-on-insulator wafers. *Japanese Journal of Applied Physics Part 2-Letters*, 41(12B):L1505–L1507, Dec 15 2002.

[60] C. Maleville, B. Aspar, T. Poumeyrol, H. Moriceau, M. Bruel, A. J. AubertonHerve, and T. Barge. Wafer bonding and h-implantation mechanisms involved in the smart-cut(r) technology. *Mat Sci Eng B-Solid*, 46(1-3):14–19, Apr 1997.

[61] J. Camassel and A. Tiberj. Strain effects in device processing of silicon-on-insulator materials. *Appl Surf Sci*, 212:742–748, May 15 2003.

[62] C. Manolatou, M. J. Khan, S. H. Fan, P. R. Villeneuve, H. A. Haus, and J. D. Joannopoulos. Coupling of modes analysis of resonant channel add-drop filters. *IEEE J. Quantum Elect.*, 35(9):1322–1331, Sep 1999.

[63] J. Brouckaert, W. Bogaerts, P. Dumon, J. Schrauwen, D. Van Thourhout, and R. Baets. Planar concave grating demultiplexer on a nanophotonic silicon-on-insulator platform. In *IEEE LEOS Annual Meeting*, pages 312–313, Montreal, Canada, October 2006 2006.

[64] C. K. Madsen and J. H. Zhao. *Optical Filter Design and Analysis*. John Wiley and Sons, Inc., 1999.

[65] K. Jinguji and M. Kawachi. Synthesis of coherent 2-port lattice-form optical delay-line circuit. *J Lightwave Technol*, 13(1):73–82, Jan 1995.

[66] J. H. McClellan, T. W. Parks, and L. R. Rabiner. A computer program for designing optimum fir linear phase digital filters. *IEEE Transactions on Audio and Electroacoustics*, AU-21(6):506, 1973.

[67] M. K. Chin and S. T. Ho. Design and modeling of waveguide-coupled single-mode microring resonators. *J Lightwave Technol*, 16(8):1433–1446, Aug 1998.

[68] A. Vorckel, M. Monster, W. Henschel, P. H. Bolivar, and H. Kurz. Asymmetrically coupled silicon-on-insulator microring resonators for compact add-drop multiplexers. *IEEE Photonic Techn. Lett.*, 15(7):921–923, Jul 2003.

[69] B. E. Little, J. P. Laine, and S. T. Chu. Surface-roughness-induced contradirectional coupling in ring and disk resonators. *Opt Lett*, 22(1):4–6, Jan 1 1997.

[70] I. Christiaens. *Vertically Coupled Microring Resonators Fabricated with Wafer Bonding*. Phd, Ghent University, 2005.

[71] H. S. Lee, C. H. Choi, B. H. O, D. G. Park, B. G. Kang, S. H. Kim, S. G. Lee, and E. H. Lee. A nonunitary transfer matrix method for practical analysis of racetrack microresonator waveguide. *IEEE Photonic Techn. Lett.*, 16(4):1086–1088, Apr 2004.

[72] M. J. Martin, T. Benyattou, R. Orobtchouk, F. Rooms, A. Morand, I. Schanen, and P. Benech. Scanning near-field optical microscopy as a tool for the characterization of multimode interference devices. *Appl Optics*, 44(13):2558–2563, May 1 2005.

[73] K. Oda, N. Takato, H. Toba, and K. Nosu. A wideband guided-wave periodic multi demultiplexer with a ring resonator for optical fdm transmission-systems. *J Lightwave Technol*, 6(6):1016–1023, Jun 1988.

[74] S. Suzuki, M. Yanagisawa, Y. Hibino, and K. Oda. High-density integrated planar lightwave circuits using sio2-geo2 wave-guides with a high refractive-index difference. *J Lightwave Technol*, 12(5):790–796, May 1994.

[75] D. Y. Wang, G. F. Jin, Y. B. Yan, and M. X. Wu. Aberration theory of arrayed waveguide grating. *J Lightwave Technol*, 19(2):279–284, Feb 2001.

[76] R. Marz and C. Cremer. On the theory of planar spectrographs. *J Lightwave Technol*, 10(12):2017–2022, Dec 1992.

[77] C. Dragone, C. A. Edwards, and R. C. Kistler. Integrated-optics n x n multiplexer on silicon. *IEEE Photonic Techn. Lett.*, 3(10):896–899, Oct 1991.

[78] M. K. Smit and C. Van Dam. Phasar-based wdm-devices: Principles, design and applications. *Selected Topics in Quantum Electronics, IEEE Journal of*, 2(2):236–250, 1996.

[79] C. Dragone. Efficiency of a periodic array with nearly ideal element pattern. *Photonics Technology Letters, IEEE*, 1(8):238–240, 1989.

[80] M. R. Amersfoort, J. B. D. Soole, H. P. LeBlanc, N. C. Andreadakis, A. Rajhel, and C. Caneau. Passband broadening of integrated arrayed waveguide filters using multimode interference couplers. *Electron Lett*, 32(5):449–451, Feb 29 1996.

[81] K. Okamoto and A. Sugita. Flat spectral response arrayed-waveguide grating multiplexer with parabolic waveguide horns. *Electron Lett*, 32(18):1661–1662, Aug 29 1996.

[82] B. Luyssaert. *Compact Planar Waveguide Spot-Size Converters in Silicon-on-Insulator*. PhD thesis, Ghent, 2005.

[83] C. R. Doerr, L. W. Stulz, R. Pafchek, and S. Shunk. Compact and low-loss manner of waveguide grating router passband flattening and demonstration in a 64-channel blocker/multiplexer. *IEEE Photonic Techn. Lett.*, 14(1):56–58, Jan 2002.

[84] Y. Barbarin, X. J. M. Leijtens, E. A. J. M. Bente, C. M. Louzao, J. R. Kooiman, and M. K. Smit. Extremely small awg demultiplexer fabricated on inp by using a double-etch process. *IEEE Photonic Techn. Lett.*, 16(11):2478–2480, Nov 2004.

[85] T. Fukazawa, F. Ohno, and T. Baba. Very compact arrayed-waveguide-grating demultiplexer using si photonic wire waveguides. *Jpn J Appl Phys 2*, 43(5B):L673–L675, May 15 2004.

[86] K. Sasaki, F. Ohno, A. Motegi, and T. Baba. Arrayed waveguide grating of 70x60 mu m(2) size based on si photonic wire waveguides. *Electron Lett*, 41(14):801–802, Jul 7 2005.

[87] F. Ohno, K. Sasaki, A. Motegi, and T. Baba. Reduction in sidelobe level in ultracompact arrayed waveguide grating demultiplexer based on si wire waveguide. *Jpn J Appl Phys 1*, 45(8A):6126–6131, Aug 2006.

[88] D. X. Dai and S. L. He. Novel ultracompact si-nanowire-based arrayed-waveguide grating with microbends. *Opt Express*, 14(12):5260–5265, Jun 2006.

[89] D. Dai, L. Liu, L. Wosinski, and S. He. Design and fabrication of ultra-small overlapped awg demultiplexer based on alpha-si nanowire waveguides. *Electron Lett*, 42(7):400–402, Mar 2006.

[90] H. Takahashi, I. Nishi, and Y. Hibino. 10 ghz spacing optical frequency-division multiplexer based on arrayed-wave-guide grating. *Electron Lett*, 28(4):380–382, Feb 13 1992.

[91] C. R. Doerr. Planar lightwave devices for wdm. In I. P. Kaminow and T. Li, editors, *Optical Fiber Telecommunications*, volume IV A, pages 405 – 476. Academic Press, 2002.

[92] P. Munoz, D. Pastor, and J. Capmany. Modeling and design of arrayed waveguide gratings. *J Lightwave Technol*, 20(4):661–674, Apr 2002.

[93] D. X. Dai and S. L. He. Accurate two-dimensional model of an arrayed-waveguide grating demultiplexer and optimal design based on the reciprocity theory. *J Opt Soc Am A*, 21(12):2392–2398, Dec 2004.

[94] A. Klekamp and R. Munzner. Calculation of imaging errors of awg. *J Lightwave Technol*, 21(9):1978–1986, Sep 2003.

[95] R. S. Jacobsen, K. N. Andersen, P. I. Borel, J. Fage-Pedersen, L. H. Frandsen, O. Hansen, M. Kristensen, A. V. Lavrinenko, G. Moulin, H. Ou, C. Peucheret, B. Zsigri, and A. Bjarklev. Strained silicon as a new electro-optic material. *Nature*, 441(7090):199–202, May 11 2006.

[96] H. Desmet, K. Neyts, and R. Baets. Modeling nematic liquid crystals in the neighborhood of edges. *J Appl Phys*, 98(12):–, Dec 15 2005.

[97] H. Desmet, W. Bogaerts, A. Adamski, J. Beeckman, K. Neyts, and R. Baets. Silicon-on-insulator optical waveguides with liquid crystal cladding for switching and tuning. In *ECOC 2003*, volume 3, pages 430–431, Rimini, Italy, 2003.

[98] L. Wooram, P. Mahn Yong, C. Seung Hyun, L. Jihyun, C. Kim, J. Geon, and K. Byoung Whi. Bidirectional wdm-pon based on gain-saturated reflective semiconductor optical amplifiers. *Photonics Technology Letters, IEEE*, 17(11):2460–2462, 2005.

[99] R. R. Patel, S. W. Bond, M. D. Pocha, M. C. Larson, H. E. Garrett, R. F. Drayton, H. E. Petersen, D. M. Krol, R. J. Deri, and M. E. Lowry. Multiwavelength parallel optical interconnects for massively parallel processing. *Selected Topics in Quantum Electronics, IEEE Journal of*, 9(2):657–666, 2003.

[100] P. Dumon, G. Morthier, T. Farrell, N. Ryan, S. Kyriakis-Bitzaros, S. Katsafouros, and V. Grivas. Estec contract 17884/03/nl/he, multi-gigabit optical backplane interconnections, final report. Technical report, 2005.

[101] P. V. Dressendorfer. Basic mechanisms for the new millenium, short course. In *IEEE Nuclear and Space Radiation Effects Conference*, volume 1998, Newport Beach, CA, July 1998 1998.

[102] R. C. Hughes. Charge-carrier transport phenomena in amorphous sio2 - direct measurement of drift mobility and lifetime. *Phys Rev Lett*, 30(26):1333–1336, 1973.

[103] Astrium. Radiation characteristics of refractive optical components. Technical Report EF OPT.RP.MF.02.75, Astrium, December 2001.

[104] J. R. Srour, C. J. Marshall, and P. W. Marshall. Review of displacement damage effects in silicon devices. *Ieee T Nucl Sci*, 50(3):653–670, Jun 2003.