UNIVERSITEIT GENT

Universiteit Gent
Faculteit Toegepaste Wetenschappen
Vakgroep Informatietechnologie

Nanofotonische Golfgeleiders en Fotonische Kristallen in Silicium-op-Isolator

Nanophotonic Waveguides and Photonic Crystals in Silicon-on-Insulator

Wim Bogaerts

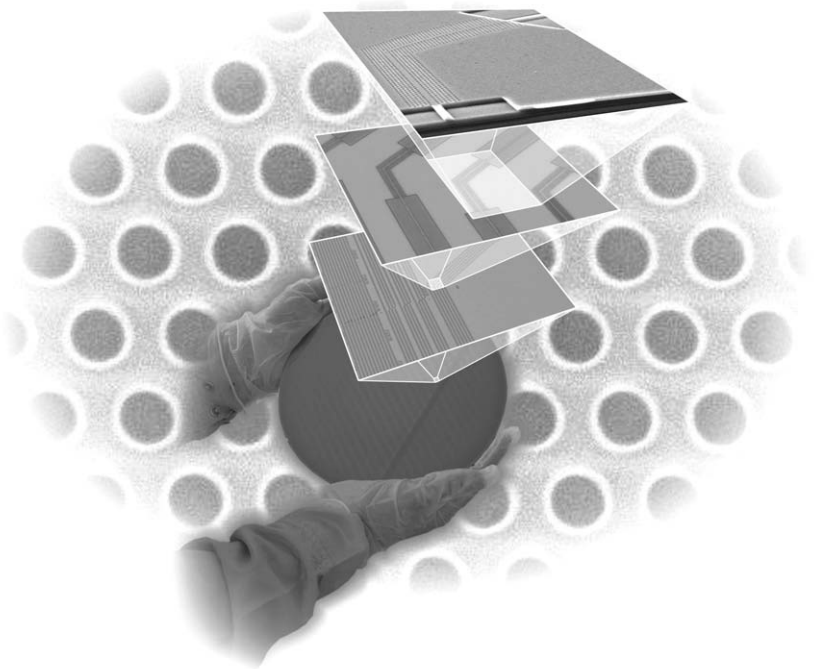

INTEC

Proefschrift tot het behalen van de graad van
Doctor in de Toegepaste Wetenschappen: Elektrotechniek
Academiejaar 2003-2004

Promotoren:

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

Examencommissie:

Prof. dr. ir. R. Verhoeven, voorzitter	Universiteit Gent, INTEC
Prof. dr. ir. D. De Zutter, secretaris	Universiteit Gent, INTEC
Prof. dr. T. Krauss	University of St. Andrews
Dr. ir. K. Ronse	IMEC vzw.
Prof. dr. I. Veretennicoff	Vrije Universiteit Brussel
Prof. dr. ir. P. Lagasse	Universiteit Gent, INTEC
Prof. dr. ir. K. Neyts	Universiteit Gent, ELIS
Prof. dr. ir. D. Van Thourhout	Universiteit Gent, INTEC
Dr. ir. P. Bienstman	Universiteit Gent, INTEC

Universiteit Gent
Faculteit Toegepaste Wetenschappen

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

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

Dit werk kwam tot stand in het kader van een specialisatiebeurs toegekend door het IWT (Vlaams Instituut voor de Bevordering van het Wetenschappelijk-Technologisch Onderzoek in de Industrie).
Part of this work was carried out in the framework of the European IST-PICCO project.

voor Katrien

Nederlandstalige Tekst

Voorwoord

Dit werk gaat over *nanofotonica*: het manipuleren van licht op een zeer kleine schaal. Licht dat informatie verspreidt over heel onze planeet. Licht voor optische vezelcommunicatie. Vijf jaar geleden beleefde de *dot.com*-rage haar gloriedagen en de ambities van de telecommunicatie-industrie kenden geen grenzen. Nanofotonica, waaronder fotonische kristallen, zouden de huidige fotonische componenten vervangen door veel compactere circuits en bovendien ongekende functionaliteit toevoegen; een proces dat zich voor de elektronica enkele decennia eerder had voltrokken. In het begin van de 21^{ste} eeuw zou jan en alleman het internet kunnen afschuimen via een persoonlijke optische vezelverbinding. Of toch niet? *Dot.com* werd snel ingeruild voor *dot.kwam* en vele ambities kwamen (tijdelijk) in de diepvries terecht. *Meer* en *beter* zijn vervangen door *goedkoper*. Gelukkig draagt nanofotonica ook deze belofte in zich, en het onderzoek is niet stilgevallen; het brandt alleen op een iets lager pitje.

Toen ik begon te werken aan mijn doctoraat was ik de eerste persoon in de groep die zich eens ging storten op het nieuwe onderzoeksgebied van fotonische kristallen. Alhoewel er al behoorlijk wat ervaring was in de groep wat betrof periodieke structuren, was het gebruik van fotonische kristallen met een hoog contrast voor golfgeleidercircuits nog onontgonnen terrein. In die tijd leek het alsof fotonische kristallen alle problemen in de fotonica zouden oplossen. Deze structuren hebben namelijk een aantal merkwaardige eigenschappen waardoor het mogelijk zou moeten zijn om zeer complexe componenten te maken bestaande uit eenvoudige bouwstenen; een beetje zoals spelen met Lego. Aan dit concept is nog niet getornd. Alleen bleek het allesbehalve eenvoudig om de meest simpele basis-Legoblokjes te maken.

In de laatste vijf jaar zijn fotonische kristallen uitgegroeid tot een belangrijke onderzoekspoot van de groep. Ten minste vijf mensen buigen zich nu dagelijks over problemen die in meer of mindere mate met fo-

tonische kristallen gerelateerd zijn. In die vijf jaar zijn we de verschillende aspecten van het optisch gedrag van deze structuren beginnen begrijpen. In dit werk proberen we die inzichten op een rijtje te zetten.

Zoals bij veel onderzoek zijn de resultaten in dit werk niet de vrucht van een eenzaat die acht uur per dag zijn uren klopt. Vele ideeën die hier zijn neergeschreven zijn voortgesproten uit de vele gesprekken met collega's, zowel in onze onderzoeksgroep als erbuiten. Alhoewel ik nu het risico loop om mensen te vergeten, wil ik hier toch een aantal mensen persoonlijk bedanken.

Bedankt, Dirk en Bert, om met jullie structuren toch eens een foton in mijn golfgeleiders te manoeuvreren, en voor de leuke momenten, vooral wanneer we 'per ongeluk' in San Fransisco terecht gekomen waren terwijl we eigenlijk tien dagen later in LA moesten zijn.

Bedankt, Pieter, Joris en Hans. Het was een plezier om jullie thesiswerk te begeleiden en een aantal van jullie resultaten zijn ook verwerkt in dit boek. Pieter, je hebt me bovendien ontlast van veel soms vervelende meet- en simulatiewerkjes op het moment dat ik er geen tijd meer voor had.

Bedankt, Peter en Lieven, voor CAMFR. De goeie 'product support' maakte het wel leuk om met vraagstukken af te komen waar de software niet op voorzien was. Zonder CAMFR zou er hier zeker een volledig hoofdstuk ontbreken.

Bedankt, Roel, omdat je het mogelijk gemaakt hebt om me op dit onderwerp te storten. De vele enthousiaste discussies, vaak doorspekt met een gezonde dosis kritische geest, hielden me op het rechte pad. En niet te vergeten de plots opduikende deadlines van de vorm 'Kunt-gedat-eens-tegen-gisteren-doen', die zeer efficiënt verhinderden dat mijn aandacht zou verslappen.

Bedankt, bureaugenootjes Bart, Freddie, Watje, Hannes en voor korte tijd Pablo. Samen hebben we duizenden elektronische wormen afgeslacht. Daarmee wil ik ook de ganse 39 bedanken: het gebouw omdat het er nog altijd staat en we dus weer eens een afbraakbarbecue kunnen houden, en de bewoners (inclusief uw dienaar) die de trapzaal soms in een (waterig) slagveld veranderden. De voltallige OCS-groep, of beter de *photonics* groep, vormden een familie van collega's die het niet alleen de moeite maakten om elke dag te gaan werken, maar om soms na de groepsvergaderingen nog enen te gaan pakken in de Vooruit of de Backstage. En we mogen de jaarlijkse *Opto-dag* natuurlijk niet vergeten.

Ik kan de kring nog ruimer maken. Vele mensen op de vakgroep INTEC doen dagdagelijks hun best om ons een hoop kopzorgen te be-

sparen, waardoor we ons kunnen concentreren op hetgeen voor ons belangrijk is (zoals de Opto-dag).

Dit werk zou vrijwel onmogelijk geweest zijn zonder de mensen van IMEC in Leuven, die er massa's siliciumplakken hebben doorgedraaid (soms met een verkeerd proces weliswaar) om onze fotonische kristallen gefabriceerd te krijgen. Dank je wel, Vincent, Johan, Stephan, Diziana, Rudi en Johan.

Ook de partners van het PICCO-project krijgen een oprecht dank je wel. De drie jaren van dit Europees project waren enorm stimulerend en veel te kort. De halfjaarlijkse projectvergaderingen waren op zich al tof om naar toe te gaan. Veel van mijn werk vond plaats in het kader van dit project en zonder de vele samenwerkingen zou dit boek er helemaal anders hebben uitgezien.

Bedankt ook aan al de mensen die minstens één fout uit dit boek gehaald hebben (en liefst voor het naar de drukker ging).

Naast mijn 'business contacts' zijn er nog andere mensen die mij overeind hebben gehouden in al die jaren. In zekere zin is dit boek ook een vrucht van al hun moeite. Mams en paps, dank je wel voor alle steun (inclusief de centjes) die me door de univ geholpen heeft. Dank je wel, kleine Goedele, omdat ik door jou 's avonds niet meer aan mijn werk hoef te denken. En bovenal, dank je wel, Katrien. Je stond altijd klaar voor mij, je doet dat nog steeds, en ik hoop dat ik dat nooit hoef te missen.

Wim Bogaerts
Gent, 12 December 2003

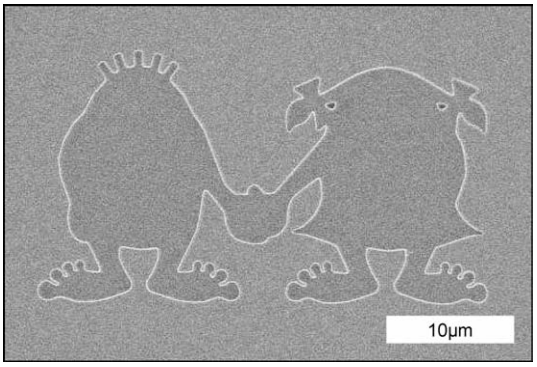

Nederlandstalige Samenvatting

1. Inleiding

1.1 Telecommunicatie en Fotonica

Zoals de titel al aangeeft, behandelen we in dit werk *nanofotonische golfgeleiders*, waaronder *fotonische kristallen*, gemaakt in *silicium-op-isolator* (SOI). Nanofotonische componenten spruiten voort uit de nood aan integratie van individuele optische componenten op een chip. *Fotonica*, of de studie van licht (fotonen) op microscopische schaal, speelt immers een belangrijke rol in de wereldwijde telecommunicatie. Optische vezels, die de ruggegraat van het internet vormen, transporteren grote hoeveelheden informatie in de vorm van pakketjes fotonen. Aan beide uiteinden van de vezel zorgen fotonische componenten voor de verwerking van de informatie en omzetting van en naar elektrische signalen. Om de capaciteit van een optische vezel ten volle te benutten, wordt de informatie verstuurd over verschillende golflengtekanalen die elk onafhankelijk door de vezel propageren. Dit noemt men *Wavelength Division Multiplexing* (WDM).

De capaciteit van de huidige telecommunicatieverbindingen wordt hoofdzakelijk begrensd door de fotonische componenten. De functionaliteit die men momenteel op een optische chip kan integreren is immers beperkt. Een belangrijke reden daarvoor is dat de huidige circuits gebaseerd zijn op zwakke golfgeleiders die het grootste deel van de beschikbare oppervlakte innemen op de component, en dit enkel om het licht van de ene functionele deelcomponent naar de andere te leiden. Nanofotonische golfgeleiders nemen een veel kleinere oppervlakte in en maken het daardoor mogelijk meer functionaliteit op

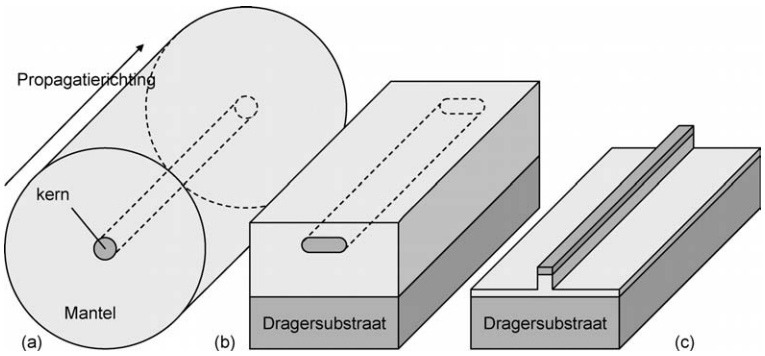

Figuur 1: Golfgeleiderdoorsnedes. (a) Een monomode optische vezel, met een $9\mu m$ dikke kern uit glas en een mantel van $125\mu m$ diameter. Het index contrast tussen kern en mantel is typisch kleiner dan 1%. (b) Een golfgeleider in *silica-op-silicium*, met gelijkaardige kernafmetingen en brekingsindexcontrast. (c) Een hoog-contrastgolfgeleider in *silicium-op-isolator* met submicron afmetingen. De kern uit silicium heeft een brekingsindex van 3.45, in contrast met een mantel met index 1.45 en lucht met een index 1.0.

eenzelfde component te integreren zonder dat daardoor de component groter wordt.

De meeste optische golfgeleiders zijn gebaseerd op het principe dat licht zich bij voorkeur voortplant in een materiaal met hoge brekingsindex. Figuur 1 toont een aantal golfgeleiders die volgens dit principe werken. De optische vezel is een voorbeeld van een zwakke golfgeleider. Het materiaal van de kern heeft een brekingsindex die slechts een fractie hoger is dan die van de omringende mantel. Het licht bevindt zich in een diffuse verdeling (een *mode*) in en rond deze kern [1]. De meeste optische vezels die voor telecommunicatie gebruikt worden ondersteunen maar één geleide mode, zogenaamde *monomode* vezels. De kern heeft typisch een diameter van $9\mu m$. Het lage brekingsindexcontrast zorgt voor een zwakke insluiting en monomode vezels zijn dan ook zeer gevoelig voor buiging. Als de vezel te sterk gekromd wordt, kan het licht als het ware 'uit de bocht vliegen'.

Voorbeeld (b) uit figuur 1 is eveneens een zwakke golfgeleider, maar ditmaal met het doel te connecties te verzorgen binnen een planair circuit, en heeft gelijkaardige afmetingen als monomode vezels. Om dezelfde reden moeten bochten in deze golfgeleiders voldoende ruim genomen worden met bochtstralen van minstens $100\mu m$. Dit verklaart waarom dergelijke structuren veel oppervlakte innemen enkel voor eenvoudige verbindingen.

1.2 Nanofotonica

Het voorbeeld in figuur 1c is een nanofotonische golfgeleider, een zogenaamde *fotonische draad*[1]. In plaats van een laag contrast tussen kern en mantel wordt hier een zeer hoog contrast gebruikt. De kern, gemaakt uit een halfgeleidermateriaal met een brekingsindex n groter dan 3, is omgeven door lucht (met een index n van 1.0) en onderaan door glas ($n = 1.45$). Om bij een dergelijk contrast slechts één geleide mode te hebben, moet een dergelijke golfgeleider voldoende kleine afmetingen hebben, van de orde van $500nm$. Door het sterke brekingsindexcontrast wordt het licht echter veel sterker ingesloten dan bij klassieke optische golfgeleiders, en zijn zeer korte bochten mogelijk.

Halfgeleiders zijn een ideaal materiaal voor nanofotonische componenten. Niet alleen hebben ze een grote brekingsindex, maar bovendien kunnen ze bewerkt worden met technologieën die ook in de elektronica gebruikt worden. Afhankelijk van de toepassingen, kan er gekozen worden voor silicium-gebaseerde materiaalsystemen, of voor de meer exotische halfgeleiders uit de $III - V$ groep van het periodiek systeem. Silicium, en dan vooral in de vorm van *silicium-op-isolator* (SOI), is een zeer goed materiaal voor nanofotonische golfgeleiders, zoals geïllustreerd in figuur 1c. De siliciumtoplaag heeft een hoge brekingsindex en is transparent voor de veel gebruikte telecommunicatiegolflengten tussen $1.3\mu m$ en $1.6\mu m$. De onderliggende laag SiO_2 is eveneens transparant maar door haar lagere brekingsindex houdt zij de toplaag, waarin de golfgeleiders gedefinieerd worden, gescheiden van het silicium dragersubstraat. $III - V$ halfgeleiders, zoals *galliumarsenide* of *indiumfosfide* maken het dan weer mogelijk om actieve functies, zoals lichtbronnen, te integreren.

Naast fotonische draden zijn ook *fotonische kristallen*[2] beloftevolle structuren om nanofotonische circuits in te ontwerpen. Terwijl fotonische draden licht insluiten door de hogere brekingsindex van de kern, wordt het licht in fotonische kristallen gecontroleerd door het samenspel van veelvuldige reflecties aan een periodiek rooster. Dit rooster kan bijvoorbeeld bestaan uit een halfgeleidermateriaal waarin periodiek gaatjes van enkele honderden nanometer geëtst zijn. Zoals bij een fotonische draad hebben we ook hier een groot brekingsindexcontrast, wat zorgt voor sterke verstrooiing van het licht aan de gaatjes. Door deze interactie van het licht met de periodieke structuur kan licht

[1] Engels: Photonic wire
[2] Engels: Photonic crystals

van bepaalde golflengtes niet meer door de structuur propageren. Er ontstaat een *fotonische verboden zone*[3]. Dit effect kan men nu gebruiken om componenten te maken: door selectief defecten in de periodieke structuur aan te brengen, wordt het licht aan deze defecten gebonden. Het kan immers niet door de omliggende structuur propageren. Hierdoor wordt een lijndefect een golfgeleider en een puntdefect een caviteit waarin het licht opgesloten kan worden. In paragraaf 2 bekijken we fotonische kristallen in meer detail.

1.3 Fabricage van Nanofotonische Componenten

Nanofotonische bouwstenen hebben typisch een kleinste afmeting van enkele honderden nanometer. De nauwkeurigheid waarmee ze moeten gefabriceerd worden is echter van de orde van $10nm$. Klassieke optische golfgeleiders kunnen nog gemakkelijk met optische lithografie gedefinieerd worden, waarbij een masker met het patroon van het circuit met een lichtbron afgebeeld wordt op een lichtgevoelige lak. Na het ontwikkelen van deze fotolak kan het patroon dan door een etsproces overgezet worden in het golfgeleidermateriaal. Dit proces is relatief eenvoudig in gebruik en is geschikt voor massaproductie. De resolutie waarmee de structuren gedefinieerd kunnen worden is echter sterk afhankelijk van de golflengte van de lichtbron. De klassieke optische lithografie, met golflengtes in het zichtbare gebied of nabije UV, schiet echter tekort voor het definiëren van nanofotonische structuren.

Voor het onderzoek naar nanofotonische golfgeleiders maakt men meestal gebruik van elektronenbundellithografie [4]. Hierbij wordt het patroon rechtstreeks in een lak geschreven met een elektronenbundel. Alhoewel met dit proces heel fijne structuren kunnen gemaakt worden, is het schrijven een tijdrovend proces. Dit maakt het proces iets minder geschikt voor massaproductie.

Het alternatief wordt aangereikt door de CMOS-industrie. Geavanceerde microprocessoren bevatten transistoren met afmetingen van de orde van $100nm$ en worden gemaakt met technieken gebaseerd op optische lithografie. De benodigde resolutie wordt verkregen door gebruik te maken van lichtbronnen in het diepe UV, bij golflengtes van $248nm$ of korter. De fabricageprocessen die rond deze technologie zijn ontwikkeld zijn echter toegespitst op het maken van micro-electronica

[3] Engels: photonic bandgap (PBG)
[4] Engels: e-beam lithography

en het is niet vanzelfsprekend om deze rechttoe rechtaan over te nemen voor nanofotonische circuits.

In dit werk hebben we het potentieel van *diep-UV-lithografie* voor de fabricage van nanofotonische golfgeleiderstructuren bestudeerd. Daartoe hebben we een samenwerking opgezet met *IMEC* in Leuven. In paragraaf 5 bespreken we de gebruikte processen in meer detail, alsook de hindernissen die we op de weg zijn tegengekomen.

1.4 Doelstellingen

Nanofotonische golfgeleiders en fotonische kristallen zijn allesbehalve eenvoudig om te modelleren, laat staan te ontwerpen of te maken. In dit werk willen we aan aantal moeilijkheden aanpakken. We hebben ons daartoe hoofdzakelijk beperkt tot 2D fotonische kristallen, vermits deze eenvoudig te integreren zijn op een planair circuit. De eerste fotonische kristalgolfgeleiders bleken echter onaanvaardbaar hoge verliezen te hebben. Om deze verliezen te beperken hebben we geprobeerd de verschillende verliesmechanismes in kaart te brengen. De verworven kennis werd dan gebruikt om betere golfgeleiders te ontwerpen.

Tegelijkertijd wilden we het tekort aan massafabricagemogelijkheden invullen. Voor de aanvang van dit werk was elektronenbundellithografie een enige mogelijkheid om nanofotonische structuren te definiëren. De toegang tot de faciliteiten van IMEC creëerde de mogelijkheid om de nodige ervaring in CMOS technologie te verwerven en een fabricageproces te ontwikkelen.

1.5 Besluit

We hebben in dit eerste deel geprobeerd de probleemsituatie te schetsen, alsook de oplossingen die we in dit werk verkend hebben. In de rest van deze samenvatting, die ruwweg dezelfde structuur heeft als de uitgebreide Engelse tekst, zullen we de volgende onderwerpen in detail behandelen. In deel 2 gaan we dieper in op de bijzondere eigenschappen van fotonische kristallen. Die kennis gebruiken we om goede golfgeleiderstructuren te ontwerpen in deel 3. Naast golfgeleiders bekijken we ook kort de in- en uitkoppelstructuren om licht van en naar optische vezels te koppelen. In deel 4 bespreken we verschillende mechanismes die bijdragen tot de propagatieverliezen in nanofotonische golfgeleiders. Deel 5 behandelt de fabricage met behulp van diep-UV-lithografie. Van de daarvoor ontworpen maskers wordt een

onverzicht gegeven in appendix A. Tot slot bespreken we in deel 6 de metingen op de gemaakte structuren.

De resultaten van dit werk zijn voorgesteld in tal van publicaties, zowel in internationale vaktijdschriften als op internationale conferenties. Een overzicht van deze publicaties vindt U in appendix B.

2. Fotonische Kristallen

In dit deel zullen we fotonische kristallen en hun eigenschappen bespreken. We spitsen ons hierbij vooral toe op 2D periodieke structuren. We zullen hierbij niet te diep ingaan op de theorie die in de Engelse tekst besproken wordt, maar ons vooral toespitsen op de eigenschappen en de toepassingen van fotonische kristallen.

2.1 Wat zijn Fotonische Kristallen?

Fotonische kristallen werden geïntroduceerd door Eli Yablonovitch in 1987 [2, 3]. Hij stelde een materiaal voor met een periodiek karakter dat opaak was voor licht binnen een bepaald golflengtegebied. Naar analogie met de halfgeleiderfysica werd dit gebied een *fotonische verboden zone* genoemd en binnen dit frequentiegebied (of golflengtegebied) kan in het materiaal geen straling propageren. Een voorwaarde voor het ontstaan van een verboden zone is dat een voldoende symmetrische kristalstructuur wordt gebruikt en dat er een voldoende groot brekingsindexcontrast is binnen de eenheidscel van het rooster [4]. De periodiciteit kan zich uitstrekken in één, twee of drie dimensies, zoals blijkt uit de voorbeelden in figuur 2.

Alhoewel in 3D-kristallen de mogelijkheid bestaat om het licht in alle richtingen volledig te controleren, is het onderzoek naar praktische toepassingen geconcentreerd op 2D-kristallen. Hierbij wordt een fotonisch kristal gebruikt om het licht in het vlak te controleren, terwijl het in de verticale richting door een gelaagde golfgeleider ingesloten wordt (figuur 3). Deze planaire fotonische kristallen[5] zijn meestal diepgeëtste halfgeleiderstructuren en er zijn reeds verschillende kristallen in het optisch golflengtegebied gedemonstreerd. Voorbeelden van dergelijke structuren zijn een driehoekig rooster van gaatjes geëtst in een GaAs-AlGaAs lagenstructuur [5, 6] of in een een SOI substraat [7, 8].

[5] Engels: Photonic Crystal Slabs

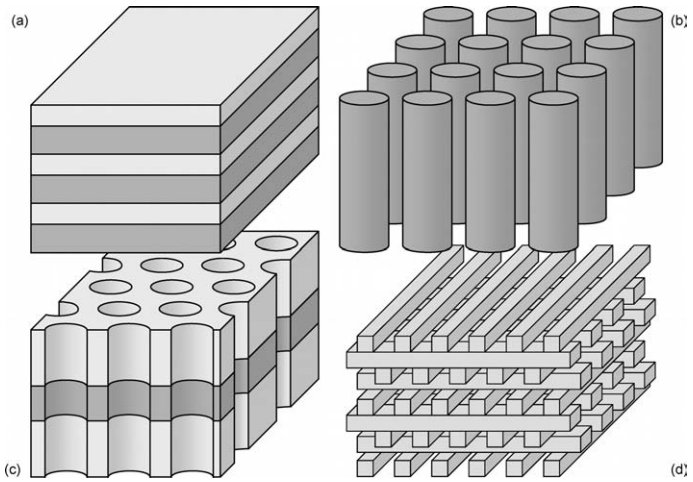

Figuur 2: Voorbeelden van fotonische kristallen. (a) Een 1D periodieke Braggspiegel. (b)Een 2D periodiek vierkant rooster van pilaartjes met hoge brekingsindex. (c) Een 2D periodieke planair fotonisch kristal met luchtgaatjes. (d) Een 3D periodieke stapeling van balken met een hoge brekingsindex.

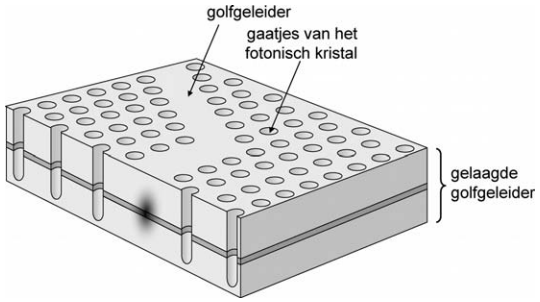

Figuur 3: Een planaire fotonische kristalgolfgeleider. In de verticale richting wordt het licht ingesloten door de hoge brekingsindex van de kernlaag. In de horizontale richting verbiedt het effect van de fotonische verboden zone dat het licht wegpropageert van het golfgeleiderdefect.

2.2 Licht en Elektromagnetische Straling

Licht is een vorm van elektromagnetische straling en wordt beschreven door de vergelijkingen van Maxwell. In een uniform medium propageert licht als een vlakke golf en de propagatieconstante k is recht evenredig met de frequentie f van het licht, ongeacht de propagatierichting. Een uniform medium schaalt ook continu; er is geen voorkeursafstand of -richting, en de elektromagnetische straling zal zich hetzelfde gedragen ongeacht de frequentie of golflengte.

Als de uniformiteit van het medium verbroken wordt, zoals bij fotonische kristallen, verandert de situatie. De structuur van het medium legt nu een voorkeursschaal op en de propagatieconstante k zal niet meer lineair afhankelijk zijn van de frequentie f. Voor sommige frequenties f zal er zelfs geen reële oplossing meer bestaan voor k, waardoor er geen propagatieconstante mogelijk is. Indien dit voor alle richtingen geldt in een bepaald frequentiegebied, spreekt men van een fotonische verboden zone (PBG).

Het verband tussen de propagatieconstante k (of bij richtingsafhankelijke systemen de propagatievector $\mathbf{k}$) en de frequentie f is de *dispersierelatie*. Het is een beschrijving van het licht in de reciproke ruimte, d.i. de fouriergetransformeerde van de reële ruimte. Het zal duidelijk worden dat een dergelijke beschrijving veel inzicht kan geven over het gedrag van licht in fotonische kristallen.

Voor een uniform medium is de dispersierelatie lineair:

$$f = \frac{k}{n}, \tag{1}$$

waarbij de evenredigheidsconstante n de brekingsindex van het materiaal is. Uitgezet in een $f|k$-diagramma zal de helling van de dispersierelatie dus kleiner zijn voor grotere brekingsindex.

2.3 2D Fotonische Kristallen

Beschouwen we de structuur uit figuur 4. Het is een driehoekig rooster van luchtgaatjes (brekingsindex $n_1 = 1$) in een achtergrondmateriaal met brekingsindex $n_2 = 3$. Als we nu de interactie van licht met een dergelijk systeem willen bestuderen, kunnen we het probleem opsplitsen in twee deelproblemen, afhankelijk van de oriëntatie van de elektromagnetische veldcomponenten. Bij de *TE-polarisatie* (transversaal elektrisch veld) zijn de elektrische veldcomponenten gericht in het vlak van het kristal, terwijl bij de *TM-polarisatie* (transversaal magnetisch

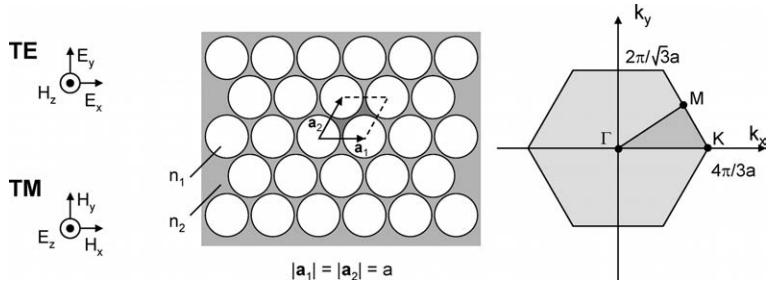

Figuur 4: Een 2D periodiek driehoekig rooster van luchtgaatjes in een materiaal met hoge brekingsindex. Vlnr.: de structuur met de eenheidscel, het driehoekig rooster en de gereduceerde brillouinzone.

veld) de elektrische veldcomponenten uit het vlak gericht zijn. Voor de magnetische veldcomponenten is de situatie net omgekeerd. Zoals we verder zullen zien is het gedrag van een 2D fotonische kristal fundamenteel verschillend voor beide polarisaties.

Beschouwen we nu het gedrag van het licht in dit rooster vanuit de k-ruimte. Voor elke periodieke structuur bestaat een *reciprook rooster*, dat de fouriergetransformeerde is van het oorspronkelijke periodieke rooster [4]. In de reciproke ruimte kan men de eerste brillouinzone definiëren rond de oorsprong. Wegens de periodiciteit van het rooster, is elk punt buiten deze eerste brillouinzone gekoppeld aan een punt binnen de eerste brillouinzone.

Het driehoekig rooster in figuur 4 heeft een ruitvormige eenheidscel met roostervectoren met grootte a. De bijhorende eerste brillouinzone is zeshoekig en eveneens afgebeeld in figuur 4. Wegens de symmetrie van het rooster kunnen we het gedrag van het rooster beschrijven in een klein deel van de brillouinzone, vermits elk ander punt door symmetrie kan afgebeeld worden op een punt in de driehoekige *gereduceerde brillouinzone*. De gereduceerde brillouinzone van het driehoekig rooster is begrensd door de punten K, M en Γ. Als we nu de dispersierelatie tussen de propagatievector $\mathbf{k}$ en de frequentie f uitzetten dan volstaat het om dit te doen langs de omtrek van de gereduceerde brillouinzone. Figuur 5 toont dit zowel voor de TE- als de TM-polarisatie. Voor de duidelijkheid wordt ook de dispersierelatie van een homogeen medium uitgezet, waarbij f recht evenredig is met de grootte van $\mathbf{k}$.

We zien dat in het fotonisch kristal voor beide polarisaties een frequentiegebied bestaat waar er geen oplossing is voor $\mathbf{k}$. Dit is de PBG. Voor de TE-polarisatie is de PBG veel groter dan voor TM. Dit heeft te

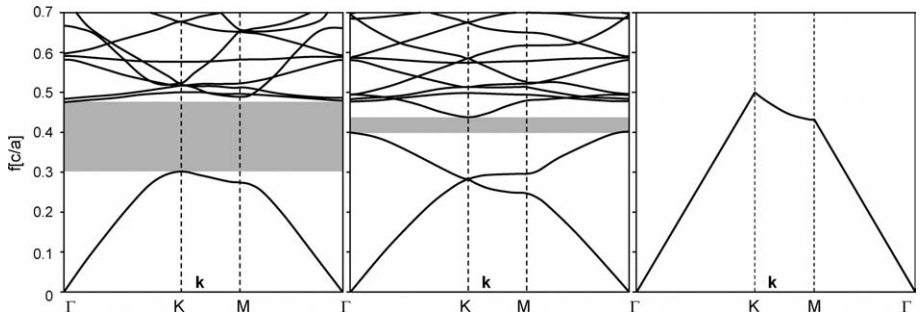

Figuur 5: Bandstructuur van het planair fotonisch kristal uit figuur 4. Vlnr.: TE-polarisatie, TM-polarisatie, een uniform medium met $n = 1.5$

maken met de geometrie van de structuur. Om een grote PBG voor TE te verkrijgen bestaat de structuur bij voorkeur uit een geconnecteerd netwerk van hoog-indexmateriaal. Voor de TM-polarisatie genieten gelokaliseerde kernen van hoog-indexmateriaal de voorkeur. Omdat het frequentiegebied opgesplitst wordt in banden met toegelaten propagatie en banden waarin geen propagatie mogelijk is noemt men een dergelijke grafiek een *bandstructuur* of *bandschema*. In dit werk, waarbij we werken met gaatjes in een hoogindexmateriaal, zullen we steeds werken met de TE-polarisatie, waarvoor de optimale kristalstructuur een driehoekig rooster van gaatjes is.

2.4 Defecten in Fotonische Kristallen

De interesse voor fotonische kristallen voor nanofotonische toepassingen is grotendeels te verklaren door de mogelijkheid om defecten aan te brengen in het periodiek rooster. Zo een defect creëert een toegelaten toestand binnen de PBG. Het licht in die toestand (of mode) is fysisch gebonden aan het defect. Het kan immers niet weglekken door het fotonische kristal wegens de PBG.

Figuur 6 illustreert een aantal defecten. Door een rij gaatjes in een fotonisch kristal te verwijderen ontstaat een lijndefect waarlangs licht kan propageren, d.i. een golfgeleider. Vermits het licht volledig ingesloten is, kan men met deze golfgeleider scherpe bochten maken zonder dat het licht kan afstralen. Op een gelijkaardige manier kan een caviteit gemaakt worden waarin het licht is opgesloten.

De dispersie in fotonische kristalgolfgeleiders verschilt sterk van die van gewone golfgeleiders. De geleide mode bestaat meestal slechts in een smal frequentiebereik en de dispersierelatie loopt zeer vlak. Dit

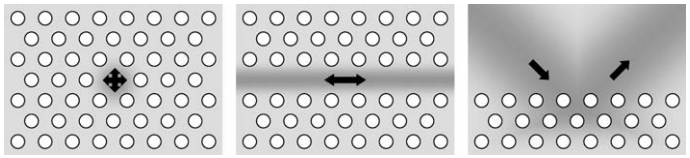

Figuur 6: Defecten in een 2D fotonisch kristal. Vlnr.: een caviteit (puntdefect), een golfgeleider (lijndefect), een spiegel (een oppervlaktedefect, dat zich voor een 2D rooster in 1 dimensie uitstrekt)

vertaalt zich naar golven die zich slechts zeer traag voortplanten door de golfgeleider. Dit effect is vooral nuttig voor toepasssingen waarbij men het licht sterk wil laten interageren met een actief materiaal.

2.5 Planaire Fotonische Kristallen en Golfgeleiders

Tot nu toe hebben we 2D fotonische kristallen besproken waarbij we veronderstelden dat de structuur zich uniform uitstrekte in de derde dimensie. In realiteit zou het licht in deze richting uitwaaieren. Om het licht nu in de verticale richting in te sluiten, wordt gebruik gemaakt van een gelaagde golfgeleider: een lagenstructuur waarbij het licht ingesloten wordt in de laag met de hoogste brekingsindex, d.i. de kernlaag, omgeven door een mantellaag. In deze gelaagde golfgeleider kan men dan een fotonisch kristal definiëren door er bijvoorbeeld gaatjes in te etsen.

Alhoewel het licht in principe ingesloten is in de verticale richting, kan de geleide mode door het fotonische kristal gekoppeld worden naar het continuüm van stralende modi in de mantellagen. Dit wordt verduidelijkt in de bandstructuur in figuur 2.25. Omdat we hier een 1D-periodieke structuur bespreken (door het defect aan te brengen is de golfgeleider nog slechts periodiek in de propagatierichting z) hebben we de bandstructuur geprojecteerd op de gereduceerde brillouinzone van de golfgeleider. We kunnen duidelijk verschillende gebieden onderscheiden. De lichtlijn begrenst het continuüm van stralingsmodi in de mantel boven en onder de structuur. De helling van de lichtlijn wordt bepaald door de brekingsindex van de mantel: hoe lager de gemiddelde brekingsindex van de mantel, hoe hoger de lichtlijn komt te liggen in het bandschema. Modes die zich boven de lichtlijn bevinden zullen steeds in meer of mindere mate afstralen. Door de projectie op de k_z-as beslaan de roostermodi van het fotonisch kristal ook een continuüm, waarin we duidelijk de opening van de PBG onderschei-

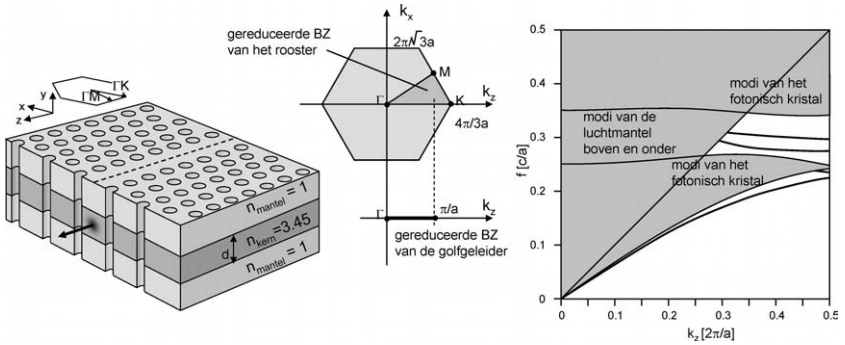

Figuur 7: Een planaire 2D fotonische kristalgolfgeleider in een siliciummembraan ($n_{kern} = 3.45$ and $n_{mantel} = 1$) en de bijhorende bandstructuur.

den. In de PBG bemerken we twee geleide modi op. In het gebied onder de lichtlijn hebben deze modi geen afstraling. Het continuüm van roostermodi bevindt zich eveneens deels onder de lichtlijn omdat het fotonisch kristal uit een hogere brekingsindex bestaat dan de luchtmantel. Nog lager vinden we opnieuw een golfgeleidermode. Die bevindt zich niet in de PBG maar is geleid omdat de golfgeleider een hogere brekingsindex heeft dan het fotonisch kristal. Er bevindt zich nu immers hoog-indexmateriaal waar er oorspronkelijk gaatjes waren.

Er spelen hier dus twee geleidingsmechanismes. Er is een zuiver *indexgeleide* mode, die zowel in horizontale als verticale richting ingesloten wordt door de hogere index van de planaire fotonische kristalgolfgeleider. We hebben ook twee *PBG-geleide* modi, die in de horizontale richting opgesloten zijn door het effect van het fotonisch kristal, maar in de verticale richting gewoon indexgeleid zijn. Boven de lichtlijn koppelen de modi met het continuüm van stralende modi en spreken we van *lekkende* modi. Merk op dat de modi door het periodiek karakter van de golfgeleider aan de rand van de brillouinzone 'teruggevouwen' worden.

In dit werk hebben we gekozen om fotonische kristalgolfgeleiders te maken voor de TE-polarisatie. Die keuze is ingegeven door het feit dat in een gelaagde configuratie de verticale insluiting beter is als de kernlaag uit een geconnecteerd netwerk bestaat (zoals een rooster van gaatjes) in plaats van uit geïsoleerde kernen (zoals een rooster van paaltjes). Omdat er dan voor de TM-polarisatie geen PBG is, wensen we die zoveel mogelijk te onderdrukken.

2.6 Modellering van Fotonische Kristallen

Voor de studie van fotonische kristallen kunnen verschillende simulatietechnieken gebruikt worden, afhankelijk van de te simuleren structuur en de gewenste data. We zullen hier kort de drie meest gebruikte technieken en hun toepassingen geven. Een eerste methode, die kan aangewend worden voor een hoeveelheid van problemen, is de *eindige differenties in de tijd* [6] (FDTD) methode [9]. Deze methode lost in essentie een gediscretiseerde benadering van de Maxwellvergelijkingen op. Zowel de ruimte als de tijd wordt gediscretiseerd en de Maxwell differentiaalvergelijkingen herleiden zich zo tot differentievergelijkingen die stap voor stap kunnen worden doorgerekend. Deze brute-kracht methode is daardoor veelzijdig en stelt weinig beperkingen aan het soort problemen die kunnen worden opgelost. Open ruimtes simuleren behoort ook tot de mogelijkheden door gebruik te maken van absorberende randvoorwaarden zoals Perfectly Matched Layers (PML)[10]. Deze gegeneraliseerde aanpak heeft echter als nadeel dat de vereiste rekenkracht het meestal onmogelijk maakt om 3D problemen volledig door te rekenen. FDTD is waarschijnlijk de meest gebruikte simulatietechniek voor fotonische kristallen.

Een tweede methode, die zeer veel wordt aangewend om de fotonische bandstructuur van fotonische kristallen te bepalen, is de ontwikkeling van het elektromagnetisch veld in vlakke golven (Plane Wave Expansion). Voor het berekenen van bandstructuren is deze methode veel sneller dan FDTD, maar ze heeft als belangrijk nadeel dat open structuren moeilijk gesimuleerd kunnen worden, omdat absorberende randvoorwaarden bij deze techniek moeilijk te realiseren zijn [11]. Toch is ook voor deze techniek een implementatie van PML mogelijk [12].

Een derde, minder gebruikte techniek, is eigenmode-expansie en -propagatie. Hierbij worden van een doorsnede van de structuur de transversale eigenmodi (zowel propagerende als stralende modi) berekend met hun propagatieconstante en het elektromagnetisch veld in de golfgeleider wordt daarna ontbonden in deze eigenmodi. Deze propageren onafhankelijk van elkaar, zolang de dwarsdoorsnede ongewijzigd blijft. Bij een overgang naar een andere golfgeleider wordt de koppeling tussen de oude en nieuwe eigenmodi berekend, zowel in reflectie als in transmissie. Om open structuren te simuleren, kunnen we ook hier PML aan de randvoorwaarden toevoegen. Deze techniek is vooral krachtig voor structuren met een constante dwarsdoorsnede of struc-

[6] Engels: Finite Difference Time Domain of FDTD

turen met veel herhaling in de propagatierichting. In INTEC hebben we twee dergelijke tools ter beschikking. FIMMPROP-3D is een commercieel pakket van Photon Design[7] waarmee 3D structuren kunnen worden uitgerekend, maar dat niet beschikt over absorberende randvoorwaarden. CAMFR[8] is een in INTEC ontwikkeld pakket dat wel met absorberende randvoorwaarden werkt.[13]. Wegens zijn efficiëntie is eigenmode-expansie veruit de meest gebruikte simulatiemethode voor de simulaties in dit werk.

2.7 Besluit

Fotonische kristallen hebben zeer interessante eigenschappen en maken het mogelijk licht te manipuleren op kleine schaal. Door deze uiteenlopende eigenschappen is het werken met fotonische kristallen echter weinig intuïtief en verre van eenvoudig. Door de kleine afmetingen en sterke indexcontrasten zijn deze structuren niet alleen moeilijk om te modelleren, maar ook de fabricage van fotonische kristallen voor het optisch golflengtegebied (inclusief de telecomgolflengten rond $1.3\mu m$ en $1.6\mu m$) is niet vanzelfsprekend.

3. Nanofotonische Componenten

In dit deel gaan we dieper in op het ontwerp van nanofotonische componenten. Hierin vormen golfgeleiders een uiterst belangrijke groep. Ze zijn niet alleen nodig voor interconnecties, maar vele functionele elementen zijn eveneens gebaseerd op golfgeleiders. Om het licht in een klein volume op te sluiten en door scherpe bochten te geleiden is er een hoog brekingsindexcontrast vereist. We bespreken hier enkele golfgeleiderontwerpen en een beperkt aantal componenten. In de uitgebreide tekst worden meer ontwerpen in detail behandeld.

Naast golfgeleidercomponenten kaarten we ook het probleem aan om nanofotonische componenten te laten interageren met de buitenwereld, in het bijzonder met optische vezels. Nanofotonische golfgeleiders hebben een veel kleinere kern dan een monomode vezel en het licht efficiënt tussen beide golfgeleiders koppelen is niet vanzelfsprekend.

[7] Photon Design website: http://www.photond.com
[8] CAMFR website: http://camfr.sourceforge.net

3.1 Materiaalkeuze: Silicium-op-Isolator

Voor planaire nanofotonische componenten hebben we een grote verscheidenheid aan (halfgeleider)materialen tot onze beschikking. Wij kozen voor *silicium-op-isolator* als basismateriaal. Het bestaat uit een dunne toplaag uit silicium die als kernlaag fungeert,met daaronder een mantel uit oxide (SiO_2) met een lagere brekingsindex. Het geheel is aangebracht op een dragersubstraat van silicium. Zoals we in het volgende deel zien, moet de oxidemantel voldoende dik zijn om de siliciumtoplaag optisch te scheiden van het substraat.

Zoals al eerder vermeld heeft dit materiaal zeer goede optische eigenschappen: zowel kern- als mantellaag zijn transparant bij telecomgolflengten en het brekingsindexcontrast tussen de kernlaag $n_{kern} = 3.45$ en de mantellaag $n_{mantel} = 1.45$ is zeer groot. Daardoor heeft de lichtlijn van de oxidemantel een grote helling, wat helpt om verstrooiingsverliezen te beperken (zie verder).

Bij het fabriceren van structuren in SOI hebben we de keuze om de structuren enkel in de toplaag of ook in de mantellaag te etsen. Door het oxide te etsen vervangen we wel een deel van de mantellaag door lucht, wat de brekingsindex van de mantel verlaagt en het optisch gedrag ten goede komt. In het geval van fotonische kristallen zijn de gaatjes bovendien dieper, waardoor het licht niet verstrooid kan worden aan de bodem van de gaatjes.

Omdat SOI een asymmetrische verticale structuur heeft, zijn de TE- en de TM-polarisaties niet volledig ontkoppeld. Voor fotonische kristalgolfgeleiders kan dit een probleem zijn omdat een fotonisch kristal met geëtste gaatjes slechts een PBG heeft voor de TE-polarisatie. Door het diepe etsen verlagen we de brekingsindex van de ondermantel en maken we de structuur dus meer symmetrisch.

Waarom zou men dan er dan nog voor opteren om de oxidemantel niet te etsen? Het antwoord, zoals we in deel 5 zullen zien, ligt bij fabricage. Het etsen van de oxidemantel introduceert immers ruwheden op de zijwanden die het licht verstrooien, wat grote propagatieverliezen met zich meebrengt.

3.2 Planaire Fotonische Kristalgolfgeleiders

Toen het concept van fotonische kristalgolfgeleiders voor het eerst geïntroduceerd werd, leek het maken van een goede golfgeleider niet veel ingewikkelder dan het wegnemen van een rij gaatjes. Het werd echter snel duidelijk dat deze zogenaamde W1 golfgeleider een allesbehalve

optimaal ontwerp was en dat het vinden van een goed ontwerp verre van eenvoudig bleek. Op de volgende pagina's bespreken we kort de vereisten voor een goede planaire fotonische kristalgolfgeleider en de redenen waarom een W1 golfgeleider daar niet aan voldoet. We geven dan ook enkele alternatieven, die in detail besproken worden in de uitgebreide Engelse tekst.

3.2.1 Vereisten voor een Goede Golfgeleider

Een goede planaire fotonische kristalgolfgeleider moet aan een aantal voorwaarden voldoen, die verduidelijkt worden in het bandschema van figuur 8.

- We wensen de TE-polarisatie te gebruiken. De koppeling met TM-modi moet dus minimaal zijn.
- We wensen een PBG-geleide mode. De PBG van het rooster moet dus voldoende breed zijn.
- We moeten werken onder de lichtlijn, om verstrooiingsverliezen tegen te gaan (zie deel 4).
- We wensen slechts 1 geleide mode.
- De mode moet voldoende bandbreedte hebben, zodat we meerdere golflengtekanalen kunnen gebruiken.
- Voor golfgeleiding wensen we geen sterke dispersie of trage golven. De helling van de dispersiekurve is dus liefst constant en voldoende steil.
- Om de golfgeleider te kunnen integreren met andere toepassingen, willen we de roosterpunten niet verplaatsen.

Bovendien moeten we rekening houden met de beperkingen van het fabricageproces. De keuze voor diep-UV-lithografie bepaalt dat de roosterperiode a onder de $400nm$ moet blijven. Eveneens kunnen we de gaten niet te groot of te klein maken. We kozen uiteindelijk voor een roosterconstante $a = 500nm$ en een gatdiameter $\oslash = 320nm$. Ook de laagdikte van de SOI is belangrijk. Onze uiteindelijke structuren werden gemaakt met een een toplaag met een dikte $d_{kern} = 220nm$ en een oxidedikte $d_{mantel} = 1\mu m$.

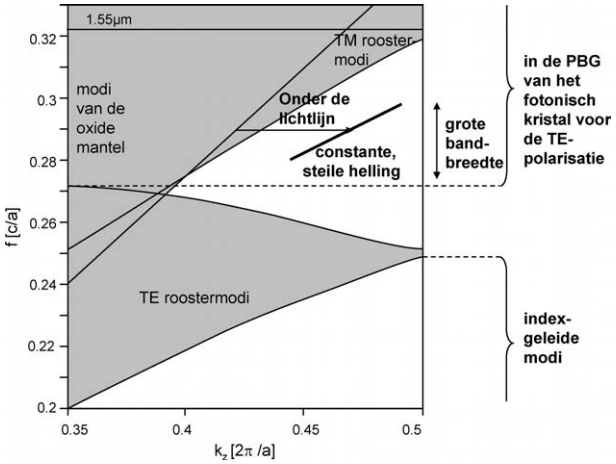

Figuur 8: Vereisten voor een SOI fotonische kristalgolfgeleider: onder de lichtlijn, in de PBG van het rooster, een constante, steile helling en voldoende bandbreedte.

3.2.2 Golfgeleiderontwerpen

Zoals al gezegd, voldoet een eenvoudige W1-golfgeleider niet aan alle vereisten die we gesteld hebben. Figuur 9 toont een diepgeëtste W1-golfgeleider en de bijhorende bandstructuur. We zien dat de TE-mode onder de lichtlijn een zeer vlakke dispersiekurve heeft. Daardoor heeft de golfgeleider een beperkte bandbreedte. Bovendien zijn er meerdere geleide modi in de PBG, waarvan een aantal boven de lichtlijn. Deze modi zullen stralingsverlies hebben. We zien ook dat de TE-mode interageert met de TM-grondmode. Dit resulteert in koppeling en het ontstaan van een zogenaamde mini-stopband (MSB), een klein golflengtegebied waar zich geen geleide mode meer bevindt.

Een W1-golfgeleider is een *indexverhogende* golfgeleider. De kern heeft immers een hogere gemiddelde brekingsindex dan het omliggende rooster. Daardoor zijn de PBG-geleide modi eigenlijk de teruggevouwen takken van de grondmode, die indexgeleid is. Deze teruggevouwen takken hebben typisch een kleine helling. Als de golfgeleiderkern een lagere gemiddelde brekingsindex zou hebben dan het omliggende kristalrooster, zou de grondmode omhoog geduwd worden tot in de PBG en krijgen we een *indexverlagende* golfgeleider.

In de uitgebreide tekst bespreken we een aantal alternatieve golfgeleiderontwerpen, geïllustreerd in figuur 10:

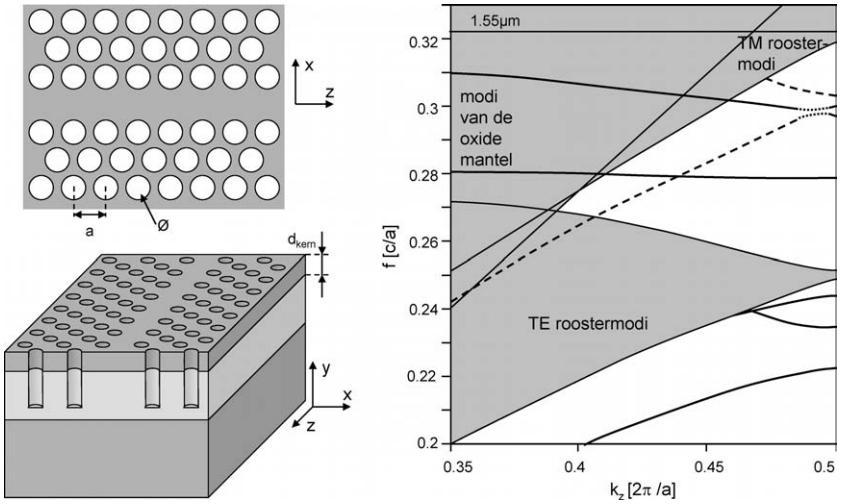

Figuur 9: Een *W1* planaire fotonische kristalgolfgeleider met bijhorende bandstructuur. Quasi-TE modi zijn aangeduid met een volle lijn, quasi-TM modi met een streepjeslijn. Modes met een gemengd karakter zijn aangeduid met stippellijn. De roosterperiode $a = 500nm$, de gat diameter $\oslash = 320nm$ en de dikte van de SOI-kern $d_{kern} = 220nm$.

- Een W0.7 golfgeleider (figuur 10a), voorgesteld voor Notomi [14], versmalt de kern tot 70% van die van een W1-golfgeleider door het rooster te verschuiven. Dit verkleint het optisch volume van de kern waardoor er slechts 1 geleide mode meer is. Alhoewel deze golfgeleider zeer lage verliezen heeft, zijn de roosterblokken verschoven ten opzichte van elkaar, wat het moeilijk maakt om functionele componenten, zoals bochten en splitsers, te integreren in deze golfgeleider.

- Een andere manier om het optisch volume van de kern te verkleinen is een deel materiaal met hoge brekingsindex wegnemen. Dit kunnen we doen door een lijndefect te creëren met kleine gaatjes (figuur 10b)2. De geleide mode van deze structuur heeft echter een oneven symmetrie, waardoor het moeilijk is om licht te koppelen naar deze mode.

- In plaats van kleinere gaatjes te gebruiken, kunnen we een defect maken met grotere gaatjes (figuur 10c). Dit is een indexverlagend defect en de grondmode van de structuur bevindt zich nu in de PBG. De mode is echter niet zeer breedbandig.

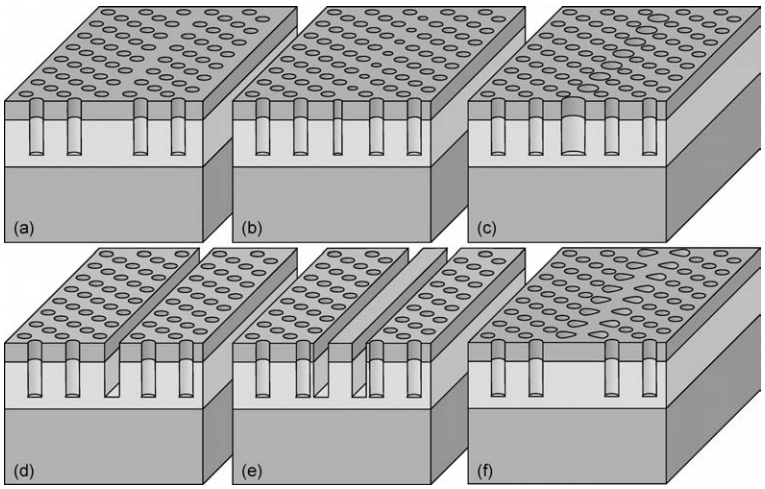

Figuur 10: Variaties op het ontwerp van een W1 golfgeleider. (a) Een W0.7 golfgeleider, waarbij het rooster verschoven is om de kern te versmallen. (b) Een W1 met een kern van kleinere gaten. (c) Een W1 met een kern van grotere gaten. (d) Een W1 met een sleuf. (e) Een fotonische draad ingebed in een fotonisch kristalrooster. (f) Een W1 met vervormde en grotere randgaten.

- Men kan ook een indexverlagend defect creëren door een sleuf te etsen in het kristal (figuur 10d). Deze golfgeleider heeft een geleide mode die symmetrisch is, maar ietwat vervormd. Bovendien is de optimale sleufbreedte van $120nm$ moeilijk te fabriceren.

- In plaats van een enkele sleuf, kunnen we een dubbele sleuf etsen (figuur 10e), waarbij we eigenlijk een fotonische draad inkapselen in een fotonisch kristal [15]. Deze golfgeleider heeft zeer goede karakteristieken en combineert de sterke punten van fotonische draden en fotonische kristalgolfgeleiders.

- In plaats van de kern te veranderen, kunnen we ook de gaatjes langs de kern veranderen (figuur 10f). Met ellips- of eivormige gaatjes verminderen we ook het optisch volume van de kern. Ook hier vinden we golfgeleidermodi met goede karakteristieken.

Bij veel van deze ontwerpen worden structuren in de kern geëtst. Alhoewel dat de karakteristiek van de golfgeleider verbetert, wordt er heel wat oppervlakte aan de zijwanden toegevoegd. Bij een imperfect etsproces kunnen ruwheden op deze zijwanden voor extra verstrooiingsverliezen zorgen.

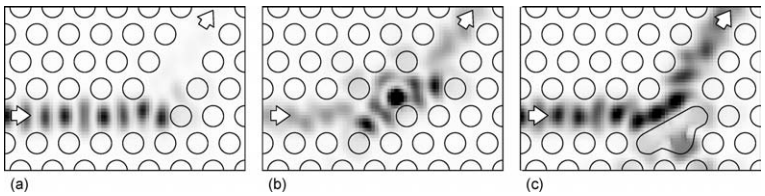

Figuur 11: Bocht met spiegel in een W1 fotonische kristalgolfgeleider. Vlnr.: De geometrie van de bocht, een veldplot bij $\lambda = 1.55\mu m$, de reflectie en transmissie als functie van de golflengte λ. De roosterperiode $a = 500nm$, de gat diameter $\oslash = 320nm$ en de dikte van de SOI-kern $d_{kern} = 220nm$.

Alhoewel elk golfgeleiderontwerp zijn sterkte kanten heeft, is er geen enkele structuur die overtuigend voldoet aan de gestelde vereisten. Het is duidelijk geen eenvoudige opgave om goede golfgeleiders in fotonische kristallen te ontwerpen.

3.2.3 Golfgeleidercomponenten

Door nu defecten toe te voegen in of bij de golfgeleider, kan men het licht sturen en golflengteselectieve functies aanbrengen. Een groot aantal componenten, zoals verschillende bochten, splitsers en caviteiten, worden besproken in de Engelse tekst. In deze samenvatting zullen we ons beperken tot een bocht van 60°. Figuur 11 toont een 2D-simulatie van drie bochtontwerpen. Alhoewel het licht in principe niet kan afstralen in het kristal, kan het wel aan de bocht gereflecteerd worden. Dit is zeer duidelijk het geval in het eerste voorbeeld, waarbij we eenvoudigweg twee W1-golfgeleiders aan elkaar gekoppeld hebben.

In het tweede ontwerp hebben we de grootte van een gat gewijzigd, waardoor er een resonantie optreedt en het licht efficiënt van de ene naar de andere golfgeleider gekoppeld wordt. Deze aanpak heeft echter enkele nadelen. In de 3D-situatie van een planair fotonisch kristal zal een dergelijke resonantie waarschijnlijk sterk koppelen naar de stralende modi van de mantel, waardoor de bocht grote verliezen uit het vlak zal hebben. Bovendien is een resonantie slechts effectief in een smal golflengtebereik, waardoor de bocht zeker niet breedbandig zal zijn.

In het derde voorbeeld hebben we geprobeerd om het optisch volume van de bocht te verkleinen door een sleuf te etsen. Hierdoor wordt het licht als het ware door de bocht gestuurd. Bovendien functioneert de wand van de sleuf als spiegel. Een dergelijk *sturend* ontwerp maakt

geen gebruik van resonanties en is daardoor breedbandiger. Het is dan ook te verkiezen boven een resonant ontwerp. Deze ontwerpsregel gaat ook op voor splitsers, of voor elke andere component die breedbandig moet zijn.

Aan de andere kant kunnen resonante caviteiten gebruikt worden om golflengteselectieve functies in te bouwen in het fotonisch kristal. De caviteiten kunnen hierbij in de golfgeleider liggen [16, 17], waardoor ze slechts bepaalde golflengtekanalen zullen doorlaten, of twee golfgeleiders met elkaar koppelen, waarbij licht bij de resonante golflengtes van de ene naar de andere golfgeleider gekoppeld zal worden. Er zijn zelfs caviteiten gedemonstreerd die het mogelijk maken licht van een bepaalde golflengte vanuit een fotonische kristalgolfgeleider te laten afstralen in de verticale richting, waar het opgevangen kan worden door een vezel of een detector [18].

Het blijkt echter niet vanzelfsprekend te zijn om fotonische kristalcaviteiten met een voldoend hoge kwaliteitsfactor Q te maken. Voor WDM-toepassingen moet de Q toch van de orde 5000 zijn, maar in de literatuur zijn zulke waarden nog niet gerapporteerd voor golfgeleidercomponenten.

3.2.4 Dispersie en Trage Golven

Zoals eerder aangehaald, hebben fotonische kristalgolfgeleiders typisch een zeer grote dispersie, wat zich uit in een vlakke dispersiekurve. Terwijl dit nadelig is voor een goede golfgeleiderwerking, kan men hiermee wel andere functionaliteiten implementeren.

Zo werd in 1998 een zogenaamde *superprima*-effect aangetoond [19]. Hierbij zal licht dat invalt op een fotonisch kristal (voor golflengten buiten de PBG) onder een andere hoek worden afgebogen afhankelijk van de golflengte. Dit effect is gelijkaardig aan dat in een prisma maar ongeveer duizend maal sterker.

Ook voor actieve functionaliteit is de sterke dispersie van fotonische kristallen een pluspunt. Het licht wordt immers door de periodieke structuur heen en weer verstrooid en propageert daardoor traag door het kristal. Deze gedistribueerde terugkoppeling vergroot de interactietijd van het licht met het materiaal. Als het materiaal winst bevat kan op deze manier een laser gemaakt worden in een fotonisch kristal, gebaseerd op hetzelfde principe als DFB-lasers [20]. Dit is ondertussen experimenteel bevestigd voor een 2D rooster [21, 22].

Ook niet-lineaire functies, zoals het opwekken van de tweede harmonische (een signaal bij de dubbele frequentie), kunnen gevoelig versterkt worden door de periodieke structuur [23]. Dit principe kan dan gebruikt worden voor het omzetten van een signaal van één golflengtekanaal naar een ander [24]. Andere niet-lineaire toepassingen, gebruik makend van het Kerr-effect, maken het mogelijk nieuwe functies, zoals een optische flip-flop, te implementeren op kleine schaal [25, 26].

3.2.5 Besluit: Fotonische Kristalgolfgeleiders

Alhoewel het principe van fotonische kristalgolfgeleiders eenvoudig is, blijkt het niet evident te zijn om goede rechte golfgeleiders te ontwerpen. Eenvoudige componenten, zoals bochten en splitsers, zijn evenmin simpel. Daarom zijn fotonische kristallen geen goede kandidaat voor interconnecties, waarbij het licht over lange afstanden op een chip getransporteerd moet worden.

Voor functionele bouwblokken zijn fotonische kristallen beter geschikt. Vooral de sterke dispersie maakt het mogelijk om actieve functies op een kleine schaal te implementeren.

3.3 Fotonische draden

Het principe van een fotonische draad verschilt niet wezenlijk van dat van een conventionele optische golfgeleider. Het licht wordt ingesloten in de kern met hoge brekingsindex. Het onderscheid wordt gemaakt door het sterke brekingsindexcontrast en de bijhorende submicron afmetingen.

De dispersierelatie van fotonische draden is veel minder exotisch dan die van fotonische kristalgolfgeleiders en dit maakt fotonische draden dan ook makkelijker in gebruik voor eenvoudige golfgeleidertoepassingen.

3.3.1 Rechte Golfgeleiders

Fotonische draden vereisen een sterke horizontale insluiting. In de verticale richting is een hoog brekingsindexcontrast geen strikte vereiste. Het is echter wel belangrijk dat een fotonische draad slechts één geleide mode heeft. Hiervoor mag de fotonische draad niet te breed zijn. Voor SOI met een kerndikte $d_{kern} = 220nm$ houden we de kernbreedte w_{kern} liefst kleiner van $550nm$.

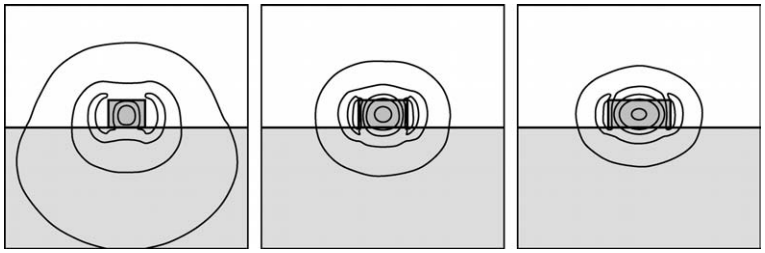

Figuur 12: Modeprofielen van de TE grondmode van fotonische draden met kernbreedte $w_{kern} = 0.3, 0.4, 0.5\mu m$. De kerndikte $d_{kern} = 0.22\mu m$, en de oxidemantel wordt oneindig dik verondersteld. Als de kern smaller wordt, strekt de mode zich verder uit in de verticale richting.

Figuur 12 geeft de modeprofielen van fotonische draden met verschillende breedtes weer. We zien dat voor smallere golfgeleiders de mode-oppervlakte toeneemt en dat de mode zich zowel horizontaal als verticaal verder uitstrekt. Vooral dat laatste heeft een belangrijke invloed op de propagatieverliezen. Voor SOI, waarbij een oxidemantel de kern van het siliciumsubstraat scheidt, mag de golfgeleidermode zich niet te ver in de mantel uitstrekken. Hierdoor zou de mode immers koppelen naar het substraat met een hoge brekingsindex en weglekken. Dit wordt besproken in het volgende deel.

3.3.2 Bochten en Spiegels

Door de sterke insluiting van het licht kunnen er scherpe bochten gemaakt worden. Dit is geïllustreerd in figuur 13, waarbij het licht met een bochtstraal van $3\mu m$ wordt omgebogen zonder noemenswaardige afstraling. Alternatief kunnen we het hoge brekingsindexcontrast gebruiken om het licht te reflecteren [27, 28, 29]. Een dergelijke hoekspiegel met een transmissie van 75% is eveneens weergegeven in figuur 13.

We kunnen op deze kleine schaal ook splitsers en koppelaars ontwerpen, zoals we uitgebreider bespreken in de Engelstalige tekst.

3.3.3 Ringresonatoren

Vermits het mogelijk is om zeer korte bochten te maken in een fotonische draad, kunnen we hiermee ook ringresonatoren ontwerpen [30]. In een ringresonator propageert het licht in een ringvormige golfgelei-

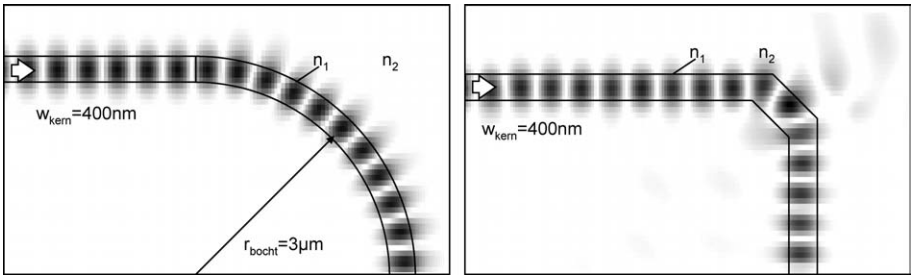

Figuur 13: 2D simulaties van een bocht in een fotonische draad. Links: een bocht met straal $r_{bocht} = 3\mu m$. Rechts: een hoekspiegel $n_1 = 2$, $n_2 = 1$, $w_{kern} = 0.4\mu m$.

der. Als de golflengte van het licht een geheel aantal maal in de omtrek 'past' (bij resonantie), dan zal het licht constructief interfereren.

We kunnen een ringresonator nu koppelen aan een golfgeleider met behulp van een directionele koppelaar, zoals voorgesteld in figuur 14. Waar de ring en de golfgeleider dicht bij elkaar komen, koppelt het licht van de één naar de ander. Door nu een ringresonator te koppelen met twee golfgeleiders zal de resonator het licht van de ene naar de andere golfgeleider koppelen bij resonantie, maar dit niet doen als de golflengte te veel afwijkt van de resonantiegolflengte. Dit is duidelijk zichtbaar in figuur 14.

De koppeling tussen de rechte golfgeleider en de ring is van cruciaal belang voor het gedrag van de ring. Daarom moet de breedte van de sleuf tussen beide golfgeleiders goed gecontroleerd kunnen worden.

3.3.3 Besluit: Fotonische Draden

Fotonische draden zijn minder complex dan fotonische kristallen voor het ontwerpen van golfgeleiderstructuren. Bovendien is het ook mogelijk om scherpe bochten te maken.

3.4 Koppelstructuren

Fotonische kristallen en fotonische draden maken het mogelijk de afmetingen van planaire golfgeleidercircuits drastisch te reduceren. De kleine golfgeleiders maken het echter moeilijk om het licht te koppelen van en naar een optische vezel. Monomode optische vezels hebben typisch een kerndiameter van $9\mu m$, terwijl onze fotonische draden in SOI een kern hebben van $0.5 \times 0.22\mu m^2$. Dit maakt, als beide golfgelei-

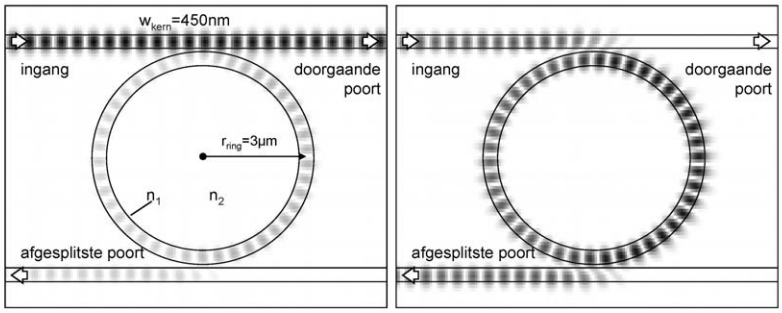

Figuur 14: Een 2D simulatie van een ringresonator in een fotonische draad. Links: niet bij de resonantiegolflengte ($\lambda = 1555nm$). Rechts: bij de resonantiegolflengte ($\lambda = 1530nm$). $w_{core} = 400nm$, $w_{sleuf} = 150nm$, $r_{ring} = 3\mu m$.

ders rechttoe rechtaan gekoppeld worden, dat de koppelefficiëntie van de orde 0.1% is.

Het is weliswaar eenvoudig om de mode van een nanofotonische golfgeleider te verbreden door een adiabatische koppelstructuur, maar het is veel moeilijker om ook de hoogte van de mode aan te passen. Een mogelijke oplossing bestaat erin om het licht naar een intermediaire golfgeleider te koppelen. Die kan bestaan uit een polymeerlaag die bovenop de nanofotonische golfgeleider is gelegd [31], of een meerlaags-halfgeleiderstructuur waarin een tweede golfgeleider zich onder de fotonische draad bevindt [32, 33].

Beide oplossingen zijn reeds met goed gevolg experimenteel aangetoond, maar het nadeel van deze aanpak is dat de koppelstructuur adiabatisch is en daardoor relatief lang, van de orde van $1mm$. Dit is in tegenspraak met de principes van nanofotonica, waar we de benodigde chipoppervlakte willen beperken.

3.4.1 Verticale Vezelkoppelaars

In plaats van een adiabatische koppelstructuur, kunnen we het licht via een 1D-periodiek oppervlakterooster koppelen naar een vezel die zich verticaal boven de golfgeleider bevindt[9]. Dit concept is geschetst in figuur 15. Het rooster is een tweede-orde rooster, met een grotere periode (typisch $580 - 600nm$) dan we normaal gebruiken voor fotonische kristallen. Ook de benodigde etsdiepte is verschillend. Afhankelijk van de ontwerpparameters, moet het rooster slechts $50nm$ tot $80nm$ diep geëtst worden. Simulaties voorspellen koppelefficiënties naar een optische vezel van 28% of meer[34].

Het gebruik van een 1D-rooster heeft zijn beperkingen. Het rooster interageert slechts met één vezelpolarisatie. Vermits in de meeste optische vezels de polarisatie kan variëren, kan het zijn dat ofwel alle licht ofwel geen licht wordt overgekoppeld. Dit kan opgelost worden door een 2D-rooster te gebruiken, zoals geïllustreerd in figuur 15. Beide polarisaties uit de vezel worden dan elk naar een eigen golfgeleider gekoppeld. De polarisaties kunnen dan in een identiek circuit verwerkt en daarna samengevoegd worden in een vezel met eenzelfde roosterstructuur [35].

[9] De ontwikkeling van deze roosterkoppelaar is het werk van ir. Dirk Taillaert, uitgevoerd in het kader van zijn doctoraatsonderzoek.

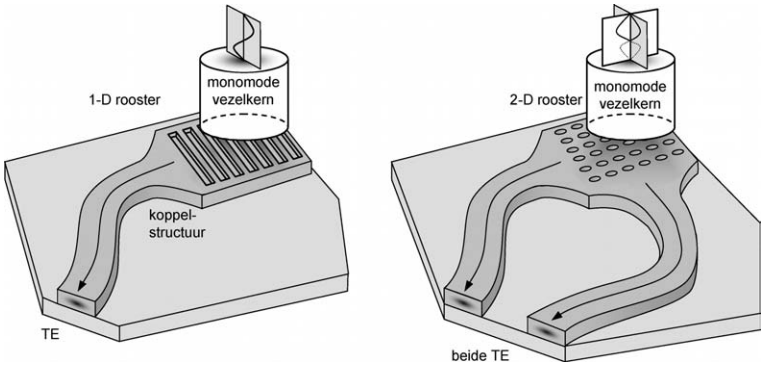

Figuur 15: Principe van een roostergebaseerde vezelkoppelaar. Licht wordt door een sterk rooster vanuit de verticale vezel in een dunne, horizontale golfgeleider gekoppeld. Links: een 1D roosterkoppelaar voor een enkele vezelpolarisatie. Rechts: een 2D roosterkoppelaar die beide vezelpolarisaties elk naar een eigen golfgeleider koppelt.

3.4.2 Compacte Koppelstructuren

De roosterkoppelaars koppelen van een vezel naar een $10\mu m$ brede golfgeleider in SOI. Er is dan een extra koppelstructuur nodig om het licht in een $500nm$ brede golfgeleider te focusseren. Een adiabatische koppelstructuur zou opnieuw veel te lang zijn.

Een alternatief is voorgesteld in figuur 16. In plaats van de grondmode van de golfgeleider gradueel te vervormen, slaan we in een *interferentiekoppelaar* bewust vele modi aan in elke sectie[10]. Met goed gekozen parameters interfereren de modi op zo'n manier dat het licht gefocust wordt in de grondmode van de fotonische draad. Het voorbeeld in figuur 16 heeft een koppelefficiëntie van 92% en dit over een lengte van slechts $10\mu m$.

3.5 Besluit

Nanofotonische golfgeleiders maken het mogelijk om zeer compacte componenten te maken. Met het hoge indexcontrast en de kleine afmetingen gaan echter een aantal ontwerpsmoeilijkheden gepaard. Fotonische kristallen hebben zeer veelzijdige dispersieve eigenschappen, maar dat maakt het nu net moeilijk om ze in te zetten als eenvoudige golfgeleiders. Daartoe zijn fotonische draden meer geschikt. Het is dan

[10]De ontwikkeling van deze interferentiekoppelaar is het werk van ir. Bert Luyssaert, uitgevoerd in het kader van zijn doctoraatsonderzoek.

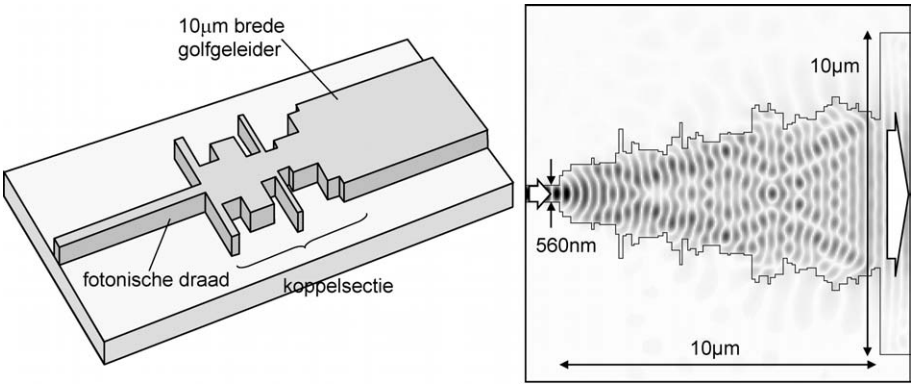

Figuur 16: Principe van een interferentie-gebaseerd koppelstructuur. Deze structuur ondersteunt een groot aantal geleide modi. De interferentie tussen al deze modi zorgt ervoor dat de grondmode van de brede golfgeleider afgebeeld wordt op de grondmode van de smalle golfgeleider. Links: Een typische structuur. Rechts: Een simulatieresultaat, met 92% transmissie.

beter fotonische kristallen in te zetten waar hun dispersie van nut kan zijn: voor actieve of golflengteselectieve toepassingen. Het is echter duidelijk dat ontwerpsregels die van toepassing zijn voor klassieke fotonische structuren niet noodzakelijk opgaan in de nanofotonica.

Om nanofotonische golfgeleiders te verbinden met de buitenwereld kunnen we roosterkoppelaars en compacte interferentiekoppelaars gebruiken. Bovendien kunnen de 2D-roosterkoppelaars gebruikt worden als polarisatiesplitser.

4. Verliezen in Nanofotonische Golfgeleiders

Nanofotonische golfgeleiders zijn weliswaar compact, maar ze bieden slechts voordeel ten opzichte van klassieke golfgeleiders indien hun propagatieverliezen voldoende laag zijn. Alhoewel fotonische draden en fotonische kristalgolfgeleiders in theorie het licht verliesloos kunnen transporteren, zijn er in de praktijk verschillende mechanismes die bijdragen tot de propagatieverliezen.

4.1 Verliesmechanismes

Fotonische kristalgolfgeleiders geleiden licht in het vlak door de PBG, terwijl in de verticale richting het licht ingesloten is door de hogere

brekingsindex van de kernlaag [6]. Dit is echter een versimpeld beeld: de golfgeleider moet eigenlijk als een 3D-systeem behandeld worden.

In een eerste benadering kan men echter stellen dat er in de geëtste gaten geen verticaal brekingsindexcontrast meer is en het licht dus niet gebonden is. Dit kan verstrooiingsverliezen veroorzaken. Zoals we verder zullen zien, heeft de keuze van de verticale lagenstructuur een zeer grote invloed op deze verliezen. Dit is een heikel punt in het onderzoek naar fotonische kristallen, vermits er een keuze moet gemaakt worden tussen III-V lagenstructuren met een laag contrast of SOI en membranen met een hoog contrast.

Ook voor verstrooiing aan ruwheden is de lagenstructuur van belang. Bij fotonische draden is dit het belangrijkste verliesmechanisme, maar ook bij fotonische kristallen moeten de ruwheden binnen de perken gehouden worden.

Een derde belangrijk verliesmechanisme is lek naar het substraat. Als de bufferlaag met lage brekingsindex niet voldoende dik is, strekt de geleide mode zich uit tot in het substraat met hoge index. Daardoor lekt het licht weg uit de golfgeleider.

4.2 Verstrooiing uit het Vlak

Het gebrek aan verticale insluiting in de geëtste gaten kan er voor zorgen dat het licht verstrooid wordt uit het vlak van de kernlaag [36]. De mate van verstrooiing is afhankelijk van het brekingsindexcontrast tussen de kernlaag en de mantellagen boven en onder. Benisty modelleerde deze verstrooiingsverliezen als stralende dipolen en bekwam dat de verliezen toenamen met het kwadraat van $\Delta\epsilon = n_{kern}^2 - n_{mantel}^2$ [36, 37]. Hij gebruikte echter storingsrekening, waardoor het model enkel geldig is voor kleine waarden van $\Delta\epsilon$ [38]. Hij behandelde de gaten eveneens als individuele verstrooiers en hield geen rekening met het periodiek karakter van fotonische kristalgolfgeleiders.

Om een beter beeld te krijgen voor grotere brekingsindexcontrasten hebben we de verstrooiingsverliezen gesimuleerd in een 1D-periodieke benadering van planaire fotonische kristalgolfgeleiders: een oneindig uitgestrekte symmetrische drielagen-golfgeleider met volledig geëtste sleuven (figuur 17). We bestudeerden daarbij het geval met een enkele luchtsleuf en het geval met een periodiek rooster van luchtsleuven. De berekeningen werden gedaan met de eigenmode-expansie-software CAMFR.

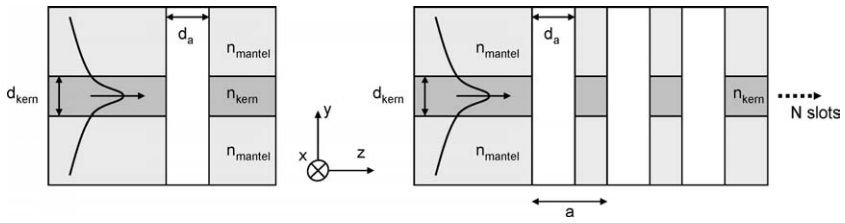

Figuur 17: Drielagen-golfgeleider met luchtsleuven. Links: een enkele sleuf. Rechts: meerdere, periodiek geschikte sleuven.

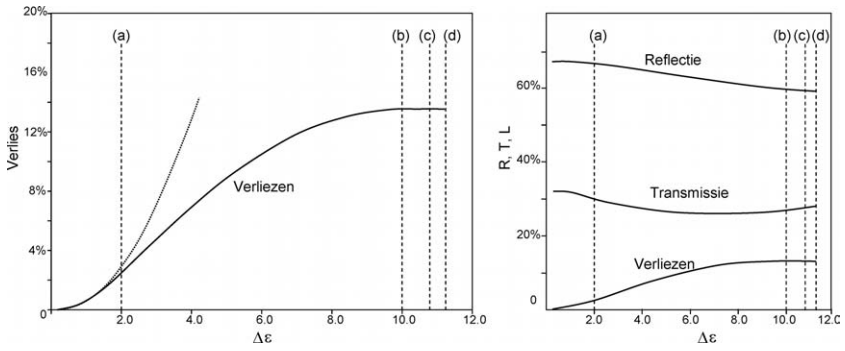

Figuur 18: Verstrooiing uit het vlak voor een enkele luchtsleuf met een breedte $d_a = 0.28\mu m$. Voor lage waarden van $\Delta\epsilon$ nemen de verliezen toe met het kwadraat van $\Delta\epsilon$ (Aangeduid met een streepjeslijn), maar ze vlakken uit voor hogere waarden. De rechtergrafiek geeft ook de reflectie en transmissie in de geleide mode.

4.2.1 Verstrooiing aan een Enkele Luchtsleuf

De verstrooiing aan een enkele luchtsleuf komt overeen met de benadering van Benisty waarbij elk gat afzonderlijk verstrooit. In figuur 18 zijn de verstrooiingsverliezen van een sleuf met breedte $d_a = 0.28\mu m$ weergegeven als functie van het brekingsindexcontrast $\Delta\epsilon$. We zien dat de verliezen inderdaad kwadratisch toenemen voor lage waarden van $\Delta\epsilon$, maar uitvlakken wanneer het brekingsindexcontrast groot wordt.

4.2.2 Verstrooiing aan Meerdere Luchtsleuven

De simulaties van een enkele luchtsleuf suggereren dat hoog verticaal brekingsindexcontrast automatisch hoge verstrooiingsverliezen geeft. Nochtans haalden we in het vorige hoofdstuk aan dat een hoog contrast net een positief effect heeft omdat de lichtlijn van de mantellaag dan

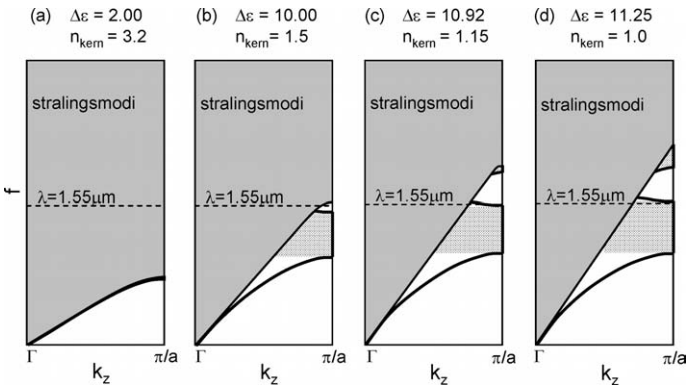

Figuur 19: Bandstructuur van planaire 1D fotonische kristallen. Het continuüm van stralingsmodi is grijs gekleurd. De fotonische verboden zone is gespikkeld. Vlnr.: a.$\Delta\epsilon = 2.0$, b.$\Delta\epsilon = 10.0$, c.$\Delta\epsilon = 10.92$, d.$\Delta\epsilon = 11.25$. De werkingsgolflengte $\lambda = 1.55\mu m$ is ook aangeduid.

een grotere helling heeft. In principe zijn de modi onder de lichtlijn dan verliesloos omdat ze niet koppelen met het continuüm van stralende modi [39, 40].

Om na te gaan welke lagenstructuur het minst gevoelig is voor verstrooiing uit het vlak, simuleerden we een periodiek rooster van luchtsleuven. We kozen de periode $a = 0.55\mu m$ en de sleufbreedte $d_a = 0.28\mu m$ zo dat we alle verschillende regimes kunnen waarnemen bij de golflengte van $1.55\mu m$ als we het verticaal brekingsindexcontrast $\Delta\epsilon$ veranderen.

Figuur 19 geeft het bandschema weer van de 1D-periodieke structuur voor de vier situaties weer die optreden als we het brekingsindexcontrast verhogen. We vinden deze situaties ook terug als we de verliezen uitzetten als functie van $\Delta\epsilon$ in figuur 20.

Voor een laag contrast (situatie a) heeft de lichtlijn een kleine helling en werken we dus steeds boven de lichtlijn. Er is dus steeds koppeling naar de stralende modi en die koppeling neemt toe als het brekingsindexcontrast toeneemt (situatie b). Op het moment dat er echter een geleide mode is onder de lichtlijn (situatie c), vallen de verliezen sterk terug. De mode is immers verliesloos en de verstrooiingsverliezen zijn nu enkel te wijten aan het overkoppelen tussen de gewone golfgeleider en de periodieke structuur. De periodieke structuur heeft ook een PBG, in deze situatie wordt het licht voornamelijk gereflecteerd (situatie d). Er is opnieuw weinig verstrooiingsverlies omdat het licht de structuur nauwelijks binnendringt.

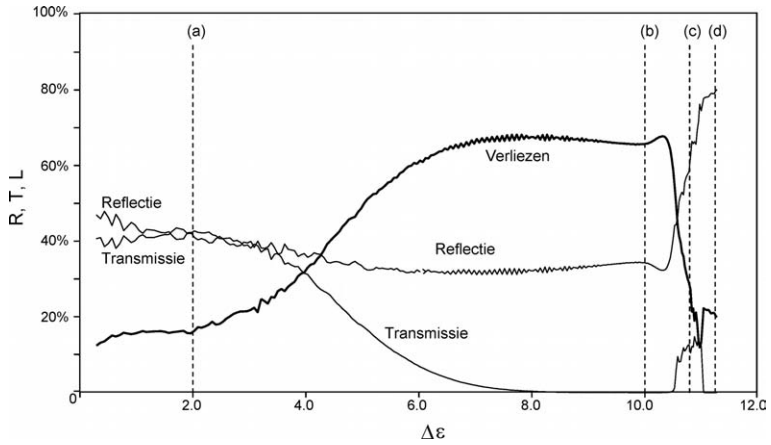

Figuur 20: Reflectie, transmissie and verstrooiingsverliezen uit het vlak voor een gelaagde golfgeleider met meerdere sleuven als functie van het verticaal brekingsindexcontrast $\Delta\epsilon$. De waarden zijn uitgemiddeld voor een structuur met 30 tot 50 periodes om de schommelingen ten gevolge van het eindige aantal periodes af te vlakken. De vier situaties uit figuur 19 zijn aangeduid.

Voor een laag brekingsindexcontrast is de verstrooiing in elk gat klein, maar naarmate het licht propageert zal er steeds een kleine hoeveelheid verstrooid worden en dit zowel in de periodieke structuur als op de plaatsen waar de periodiciteit verbroken wordt. Als het brekingsindexcontrast toeneemt, neemt ook de verstrooiing per gat toe, zodat we in een situatie komen waarbij het licht vrijwel volledig verloren gaat in de eerste periodes. Bij een geleide mode onder de lichtlijn is de periodieke structuur in principe verliesloos, zolang de periodiciteit niet verbroken wordt.

4.2.3 Besluit voor 2D-Periodieke Structuren

Uit de resultaten van de 1D-simulaties kunnen we conclusies trekken voor planaire 2D fotonische kristalgolfgeleiders [41, 38]. Voor lage contrasten zal de golfgeleider weinig verlies hebben per periode, maar accumuleren de verliezen naarmate het licht verder propageert. In bochten of aan structuren waar de periodiciteit verbroken wordt, zullen er echter geen noemenswaardig grotere verliezen optreden dan in de rechte golfgeleider.

Voor een hoog verticaal brekingsindexcontrast kan men een golfgeleider ontwerpen met mode onder de lichtlijn. Zolang de periodiciteit niet wordt verbroken is de golfgeleider verliesloos, maar elke symme-

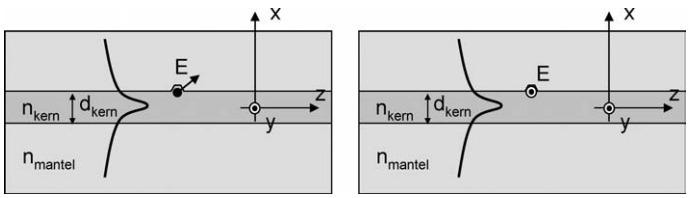

Figuur 21: Verstrooiing aan ruwheden op de zijwand van een fotonische draad (bovenaanzicht). De ruwheden worden gemodelleerd als stralende dipolen op de zijwand. Links: voor de TE mode van de golfgeleider (**H** uit het vlak). Rechts: voor de TM mode van de golfgeleider (**E** uit het vlak).

triebreuk, zoals een bocht of een splitser moet zeer zorgvuldig ontworpen worden om de verticale verstrooiingsverliezen te beperken.

Een laag verticaal brekingsindexcontrast is dus te verkiezen in de situatie waar men veel functionele elementen dicht bij elkaar integreert. In de situatie waar men lange golfgeleiders heeft met slechts sporadisch een symmetriebreuk, is een hoog brekingsindexcontrast te verkiezen, omdat men in een dergelijk systeem een golfgeleider met een geleide blochmode kan ontwerpen.

4.3 Ruwheden in Fotonische Draden

Terwijl sommige fotonische kristallen intrinsieke verstrooiingsverliezen kunnen hebben door koppeling met het continuüm van stralende modi, zijn fotonische draden in principe verliesloos zolang de dwarsdoorsnede niet verandert langs de propagatierichting. In gefabriceerde structuren zullen de geëtste zijwanden steeds een zekere ruwheid hebben. Het licht kan verstrooid worden aan deze onregelmatigheden. Deze verstrooiingsverliezen zullen ook afhankelijk zijn van het brekingsindexcontrast van de fotonische draad en van de gebruikte polarisatie. Afhankelijk van de bron, wordt gesteld dat de verstrooiingsverliezen aan ruwheden toenemen evenredig met Δn^3 [42, 43], $\Delta\epsilon$ [44, 45] of $\Delta\epsilon^2$ [46], met $\Delta\epsilon = n_{kern}^2 - n_{mantel}^2$.

Idealiter worden de verstrooiingsverliezen rigoureus in 3D gesimuleerd. We moeten ons echter beperken tot 2D simulaties in CAMFR. Figuur 21 schetst de gesimuleerde structuur: een fotonische draad in bovenaanzicht. De ruwheid op de wand modelleren we als een stralende dipool waarvan de sterkte bepaald wordt door het veld van de geleide mode op de wand. Het vermogen P dat de dipool afstraalt is

dan

$$P = \eta^2 \frac{E^2}{Z_{str}}, \tag{2}$$

met E het elektrisch veld op de wand en Z_{str} de stralingsimpedantie van de omgeving. De evenredigheidsfactor η^2 wordt bepaald door het brekingsindexverschil en het type van ruwheid. Vermits ruwheden op geëtste zijwanden hoofdzakelijk langwerpig zijn in de verticale richting, simuleerden we de verstrooiing aan een oneindig uitgestrekte ruwheid op een scheidingswand tussen twee materialen. Het horizontale brekingsindexcontrast op de scheidingswand is $\Delta\epsilon_h$. We bepaalden deze factor voor onregelmatigheden met verschillende geometrieën en vonden dat

$$\eta^2 = \gamma_{\|} \left(\Delta\epsilon_h\right)^2 \tag{3}$$

voor de situatie waarbij het E-veld parallel is met de as van de ruwheid en gelijkaardig,

$$\eta^2 = \gamma_{\perp} \left(\Delta\epsilon_h\right)^2 \tag{4}$$

voor het E-veld loodrecht op de ruwheid. Deze relatie komt overeen met de resultaten uit de literatuur [46]. We vinden bovendien dat

$$\frac{\gamma_{\|}}{\gamma_{\perp}} \approx 9. \tag{5}$$

Voor de TE-polarisatie staat het E-veld loodrecht op de verticale ruwheid en voor de TM polarisatie is het E-veld evenwijdig met de ruwheid.

Een deel van het uitgestraalde vermogen zal echter terug door de golfgeleiderkern opgevangen worden. We noemen L de fractie die verloren gaat, d.i. niet teruggekoppeld wordt naar de voorwaartse mode. L zal eveneens afhangen van de polarisatie, de brekingsindex van de mantel en de breedte van de golfgeleider.

Deze simulaties stelt ons nu in staat een maat L_{tot} voor het verstrooiingsverlies te bepalen. We kunnen weliswaar de absolute propagatieverliezen niet uitrekenen, maar we kunnen wel verschillende brekingsindexcontrasten en polarisaties vergelijken. We krijgen de volgende uitdrukkingen voor L_{tot}, afhankelijk van de polarisatie:

$$TE - mode: \quad L_{tot,TE} = \gamma_{\perp} \left(\Delta\epsilon_h\right)^2 L_{TE} \frac{E_x^2}{Z_{str,TE}}, \tag{6}$$

$$TM - mode: \quad L_{tot,TM} = \gamma_{\|} \left(\Delta\epsilon_h\right)^2 L_{TM} \frac{E_y^2}{Z_{str,TM}}. \tag{7}$$

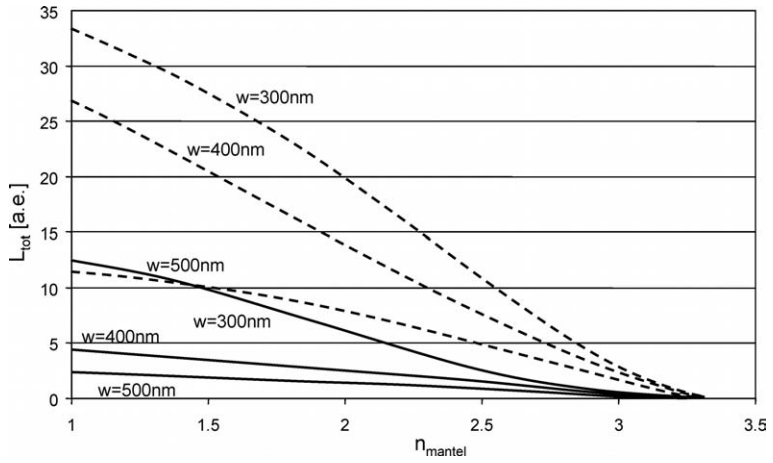

Figuur 22: Een maat voor de verstrooiing aan ruwheden op de zijwand van een fotonische draad. L_{tot}, een maat voor de verliezen uitgedrukt in arbitraire eenheden, toont de verliezen voor drie verschillende golfgeleiderbreedtes en beide polarisaties als functie van het brekingsindexcontrast. TE is is weergegeven met volle lijn, TM met streepjeslijn.

In figuur 22 is L_{tot} uitgezet als functie van de mantelindex n_{mantel} voor beide polarisaties en verschillende golfgeleiderbreedtes. We zien dat de verstrooiingsverliezen sterk toenemen met het brekingsindexcontrast. Bij een laag brekingsindexcontrast kan men de golfgeleiders echter moeilijk beschouwen als fotonische draden. De belangrijkste conclusie is hier wel dat de verstrooiingsverliezen aan ruwheden opmerkelijk groter zijn voor de TM-polarisatie dan voor de TE-polarisatie. Bovendien nemen de verliezen toe voor smallere golfgeleiders en dit omdat de veldsterkte op de wand toeneemt.

4.4 Ruwheden in Fotonische Kristallen

In planaire fotonische kristallen is verstrooiing aan ruwheden op de zijwand eveneens een bron van propagatieverliezen. Net als bij fotonische draden kunnen we het effect van het brekingsindexcontrast bestuderen. We beschouwen hiervoor opnieuw de situatie met een enkele luchtsleuf. Nu kunnen we wel het verticaal contrast expliciet in rekening brengen. De onregelmatigheid modelleren we opnieuw als een dipool, zoals afgebeeld in figuur 23.

Het is intuïtief niet onmiddellijk duidelijk of de verstrooiingsverliezen bij een hoog dan wel een laag brekingsindexcontrast groter zullen

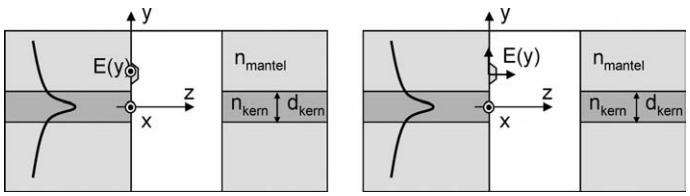

Figuur 23: Verstrooiing uit het veld aan ruwheden op de wand van een luchtsleuf in een gelaagde golfgeleider. De ruwheden worden gemodelleerd als stralende dipolen. Links: TE polarisatie. Rechts: TM polarisatie.

zijn. Bij een hoog contrast bevat de geleide mode componenten die onder een grote hoek naar boven en onder gericht zijn, maar heeft de gelaagde golfgeleider een grote numerieke apertuur die verstrooid licht makkelijk kan recupereren. Bij een laag verticaal contrast gedraagt het licht zich meer als een vlakke golf, maar de mode is meer uitgestrekt en interageert dus met een groter oppervlak.

Opnieuw gebruikten we CAMFR om de straling van de dipool te modelleren, waarbij we de dipool op de zijwand verplaatsten in de y-richting en telkens de lokale veldsterkte gebruikten om het vermogen van de dipool te berekenen. Daarbij berekenden we ook de fractie van het uitgestraalde licht dat opnieuw door de golfgeleider kon worden opgevangen, ditmaal in voorwaartse en achterwaartse richting. In een fotonisch kristal zijn beide richtingen immers gekoppeld. We hielden voor de verschillende polarisaties ook rekening met de oriëntatie van het elektrisch veld, waarbij we opnieuw veronderstelden dat de ruwheden, veroorzaakt door etsen, verticaal geöriënteerd zijn.

Dit geeft ons opnieuw een maat voor de verliezen L_{tot}, die weergegeven is in figuur 24. We zien dat de verliezen toenemen voor een lager brekingsindexcontrast. Dit omdat een golfgeleider met een hoog contrast meer van het verstrooide licht kan recupereren. De TM-polarisatie heeft bovendien duidelijk grotere verliezen dan de TE-polarisatie.

Zoals we eerder gezien hebben, hebben fotonische kristallen twee geschikte regimes om de intrinsieke verstrooiingsverliezen laag te houden. Voor verstrooiing aan ruwheden blijkt een hoog brekingsindexcontrast, zoals bij SOI of membranen, echter tot veel minder verliezen te leiden. Dit zou kunnen verklaren waarom fotonische kristalgolfgeleiders in SOI bij experimenten nog steeds veel minder verliezen hebben dan in III-V halfgeleiders [47, 48].

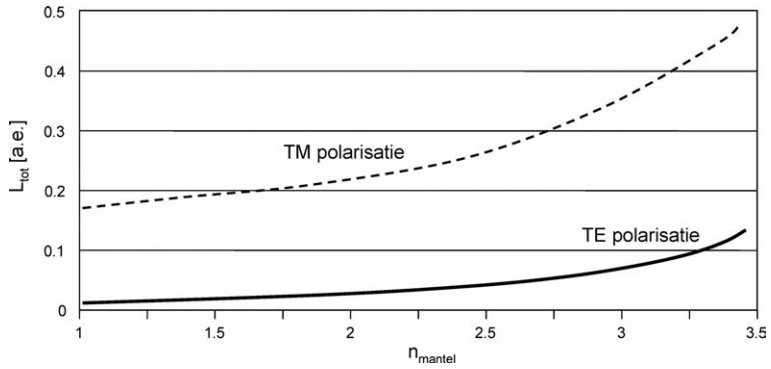

Figuur 24: Gemiddeld vermogen verloren door verstrooiing aan ruwheden op de zijwand van een enkele luchtsleuf als functie van de mantelindex n_{mantel}. De verliezen voor de TE polarisatie zijn in volle lijn aangeduid, die voor TM in streepjeslijn.

4.5 Lek naar het Substraat

Tot hier toe modelleerden we steeds drielagen-golfgeleiders. In vele reële situaties rust de gelaagde golfgeleider op een dragersubstraat dat dikwijls een hoge brekingsindex heeft. Dit is geïllustreerd in figuur 25. Als de bufferlaag te dun is, kan de geleide mode in de kernlaag het substraat 'voelen' en zal er vermogen weglekken uit de kern.

4.5.1 De Gelaagde Golfgeleider

Om een idee te krijgen van de grootte-orde van de substraatlek in SOI, gebruikten we CAMFR om de verliezen van de geleide mode te berekenen. Dit is mogelijk omdat het simulatieprogramma beschikt over PML randvoorwaarden, die het oneindig uitgestrekt substraat benaderen.

Figuur 26 toont de substraatlek in dB/mm als functie van de bufferdikte d_{mantel} voor een SOI lagenstructuur met verschillende kerndiktes d_{kern}. We zien dat de verliezen ten gevolge van substraatlek exponentieel afnemen met de bufferdikte. We zien tevens dat de verliezen voor de TM-mode veel groter zijn. Door de randvoorwaarden voor TM strekt de mode zich immers veel verder in de buffer uit. De verliezen zijn ook lager voor een dikkere kernlaag, maar om een monomode golfgeleider te hebben moet de kerndikte kleiner gehouden worden dan $255nm$.

We kozen op basis van deze simulaties een kerndikte van $220nm$. We concludeerden hieruit dat een bufferdikte van $1\mu m$ voldoende is

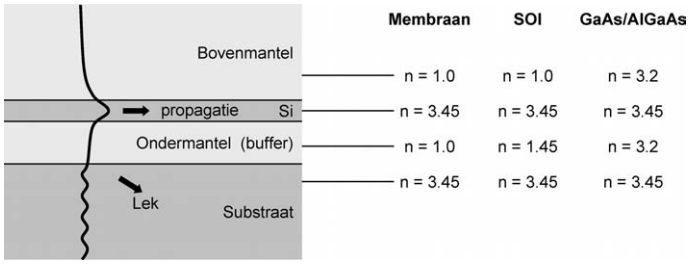

Figuur 25: Verliezen dor substraatlek. Als het substraat eenzelfde of hogere brekingsindex heeft dan de golfgeleiderkern, kan licht door de bufferlaag naar het substraat lekken.

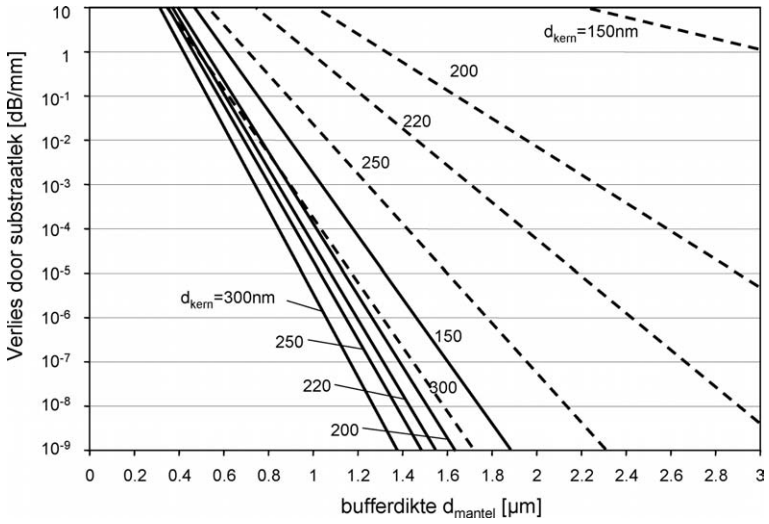

Figuur 26: Verlies door substraatlek voor een gelaagde golfgeleider als functie van de bufferdikte d_{mantel}. Het resultaat is ook weergegeven voor verschillende kerndiktes d_{kern} van de golfgeleider. De verliezen voor de TE polarisatie zijn in volle lijn aangeduid, die voor TM in streepjeslijn.

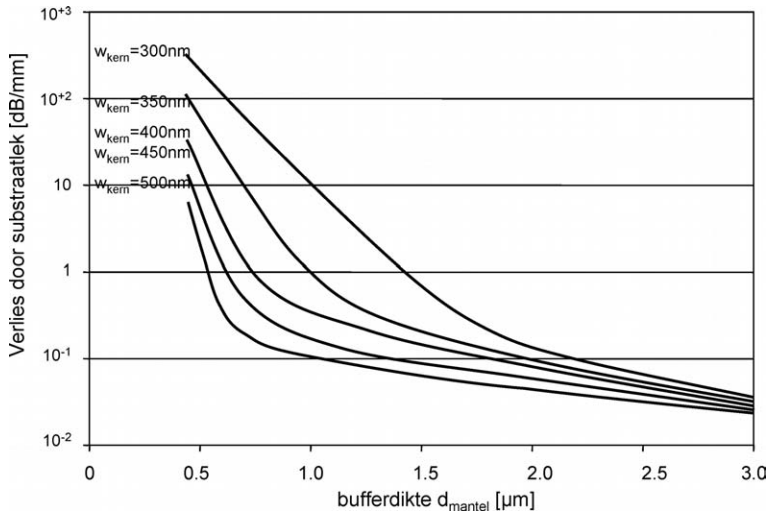

Figuur 27: Substraatlek voor fotonische draden als functie van de bufferdikte d_{mantel} voor verschillende golfgeleiderbreedtes w_{kern}.

om de TE-mode van het substraat te isoleren. Voor de TM-mode zou een bufferdikte van minstens $2\mu m$ nodig zijn.

4.5.2 Fotonische draden

De substraatlek voor een gelaagde golfgeleider geeft ons een ondergrens voor de verliezen van een fotonische draad. We hebben immers gezien dat de mode van een fotonische draad zich verticaal verder uitstrekt naarmate de golfgeleider smaller wordt.

In figuur 27 tonen we de simulatieresultaten voor de substraatlek van fotonische draden als functie van bufferdikte d_{mantel} en dit voor verschillende kernbreedtes w_{kern}. De dikte van de kern d_{kern} werd op $220nm$ gehouden. Voor een smalle golfgeleiderkern nemen de verliezen zeer snel toe.

4.6 Andere Verliesmechanismes

Naast de hier besproken verliesmechanismes kunnen we voor planaire fotonische kristalgolfgeleiders nog een aantal andere bronnen van verlies identificeren:

- Naast ruwheden op de zijwand kan het licht in een fotonisch kristal ook verstrooid worden aan de bodem van de geëtste gaatjes. Daarom moeten de gaatjes voldoende diep geëtst worden.

- In structuren met een laag verticaal contrast zal de niet-geëtste mantel onder de bodem van de gaatjes naar alle waarschijnlijkheid een grotere gemiddelde brekingsindex hebben dan de kern. Dit is een extra motivatie om de gaatjes voldoende diep te etsen.

- Als de zijwanden niet voldoende verticaal geëtst zijn, maar onderaan smaller worden, wordt het licht opnieuw in de richting van het substraat getrokken.

- In een asymmetrische lagenstructuur zijn TE en TM niet langer ontkoppeld. Waar modi met gelijke symmetrie elkaar snijden in het bandschema zullen ze interageren en kan een mini-stop-band ontstaan. Bovendien is het licht dat naar de TM-mode gekoppeld wordt niet meer ingesloten door de PBG en kan het dus afstralen.

- Zoals we in het vorige hoofdstuk hebben gezien, hebben veel planaire fotonische kristalgolfgeleiders een smalbandige geleide mode. Als nu door fabricage de parameters van de golfgeleider langs de propagatierichting licht fluctueren, zal die smalle band omhoog en omlaag schuiven in het bandschema. De geleide mode zal bij de werkingsgolflengte dan eens onder, dan weer boven de lichtlijn terechtkomen, wat sterke verstrooiingsverliezen kan veroorzaken.

4.7 Besluit

We hebben verschillende verliesmechanismes van fotonische draden als planaire fotonische kristalgolfgeleiders onderzocht. Daarbij stelden we ons de vraag welk materiaalsysteem optimaal is om de verliezen te beperken. Terwijl voor intrinsieke verstrooiingsverliezen zowel een laag brekingsindexcontrast (zoals bij III-V halfgeleiders) als een hoog brekingsindexcontrast met een geleide blochmode (zoals in SOI) geschikt is, blijkt een laag brekingsindexcontrast veel gevoeliger voor imperfecties in de fabricage. Het is ook duidelijk geworden dat de bufferlaag voldoende dik moet zijn om lek naar het substraat te vermijden.

5. Fabricage met CMOS Technologie

Nanofotonische structuren moeten met zeer grote nauwkeurigheid gemaakt worden. Bovendien hebben we gezien dat de toleranties op ruwheden op de zijwand bijzonder klein zijn. In dit deel bespreken we het fabricageproces op basis van diep-UV-lithografie voor nanofotonische golfgeleidercircuits, wat het op termijn mogelijk moet maken om componenten op commerciële schaal te fabriceren. Elektronenbundellithografie, dat meestal voor onderzoeksdoeleinden gebruikt wordt, is daarvoor niet geschikt.

Het overnemen van CMOS-processen voor het maken van nanofotonische golfgeleiders is niet vanzelfsprekend. Voor micro-elektronica wordt voor elk type structuur een geoptimaliseerd proces gebruikt. Dit is niet mogelijk bij nanofotonische componenten. Daar moeten alle deelstructuren zo nauwkeurig uitgelijnd zijn dat ze in één processtap moeten gedefinieerd worden. Nanofotonische componenten zijn ook zeer gevoelig aan wijzigingen in afmetingen. Denk aan de kleine bandbreedte van sommige fotonische kristalgolfgeleiders of de sleuf tussen een rechte golfgeleider en een ringresonator. Daarom kunnen we niet de bij CMOS gebruikelijke 10% afwijking tolereren, maar stellen we de strengere eis van 5% maximaal toegelaten afwijking. Het feit dat nanofotonische structuren zo sterk verschillen van CMOS-structuren maakt dat er veel werk in procesontwikkeling is gekropen.

5.1 Fabricageproces

Het fabricageproces bestaat uit een combinatie van lithografie en etsen. We starten hierbij van een min of meer standaard proces [49]. Eerst wordt een lichtgevoelige lak op de SOI-plak aangebracht. We gebruiken hiervoor Shipley UV3, een standaard-fotolak voor belichting bij $248nm$. Alvorens belicht te worden, wordt deze fotolak voorgebakken, waarna er een antireflectieve film op aangebracht wordt om staande-golfpatronen tijdens de belichting te voorkomen. Na de belichting wordt de fotolak opnieuw gebakken en daarna ontwikkeld. Daarna is de fotolak klaar om geëtst te worden. Het etsen van de SOI lagenstructuur gebeurt in twee stappen. Eerst etsen we de siliciumtoplaag en daarna optioneel de oxidebuffer. Dit wordt gedaan zonder de structuur tussenin aan lucht bloot te stellen. Nadien wordt de resterende fotolak verwijderd.

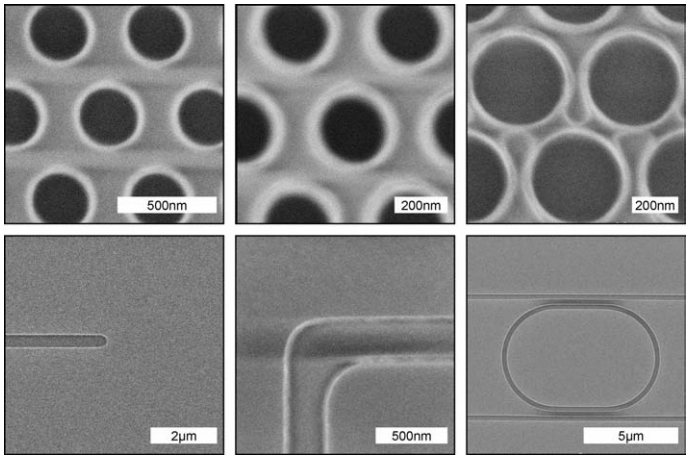

Figuur 28: Patronen in de fotolak van nanofotonische structuren, gemaakt met een elektronenmicroscoop. Bovenste rij: $300nm$ gaatjes met een periode van $500nm$, $240nm$ gaatjes met een periode van $400nm$, $420nm$ gaatjes met een periode van $500nm$. In het laatste geval zien we dat de fotolak tussen de gaatjes gedeeltelijk belicht is. Onderste rij: fotonische draden een een resonator.

5.2 Diep-UV-Lithografie

Voor diep-UV-lithografie maakten we gebruik van de technologie van IMEC in Leuven. Voor de CMOS-onderzoeksactiviteiten beschikt men daar over diep-UV-*belichters* met een golflengte van $248nm$ en $193nm$.

5.2.1 Invloed van de Belichtingsdosis

De belichters beelden eenzelfde patroon meermaals af op een $200mm$ SOI-plak en kunnen daarbij andere belichtingcondities gebruiken bij elke stap. Dit maakt het mogelijk om op relatief eenvoudige wijze een proces te karakteriseren en voor verschillende structuren de optimale procesparameters bepalen.

Figuur 28 toont enkele beelden van nanofotonische structuren in fotolak. Door de belichtingsdosis te verhogen kunnen we nu de diameter van de gaatjes doen toenemen. Voor de fotonische draden, of alleenstaande lijntjes, zal de lijnbreedte afnemen bij toenemende belichtingsdosis. Dit is geïllustreerd in figuur 29.

We zien dat we door het veranderen van de belichtingsdosis een ruime waaier aan gatgroottes en lijnbreedtes kunnen bestrijken. Voor de gaten waarbij de ontwerpdiameter 60% van de periode is, zien we bovendien dat de diameter niet te snel toeneemt. Dit geeft ons een

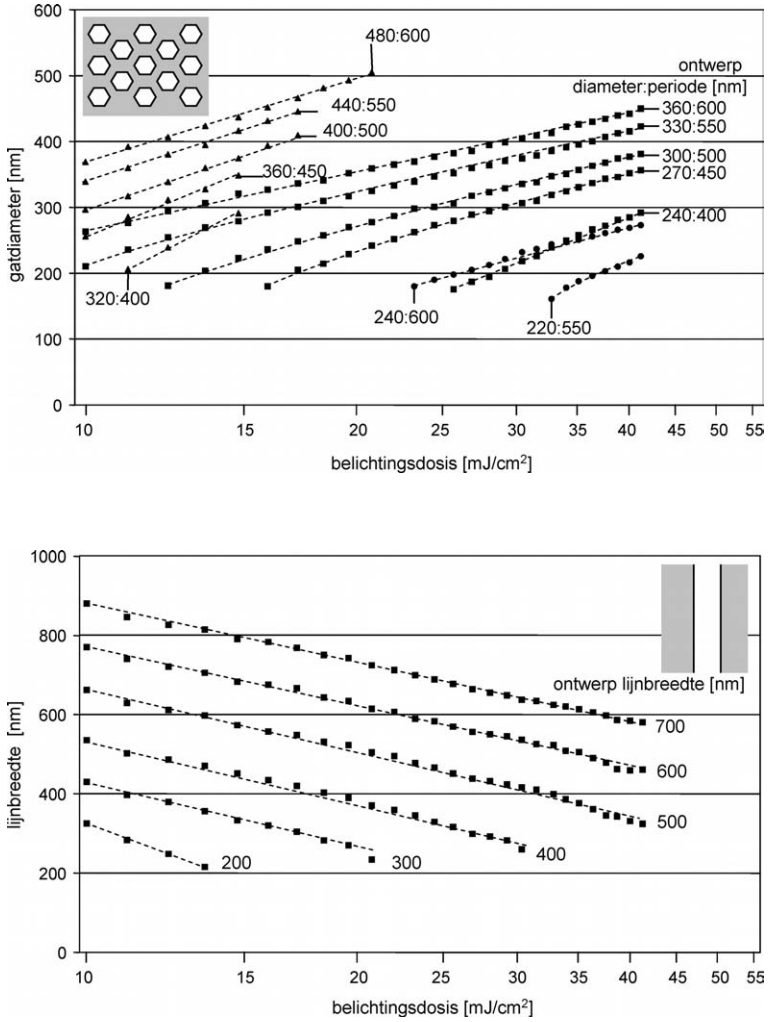

Figuur 29: Invloed van de belichtingsdosis op de afmeting van structuren. Bovenaan: diameter van de gaatjes in een driehoekig rooster als functie van de belichtingsenergie voor verschillende ontwerpkeuzes van periode en vulfactor. De verhouding diameter:periode bij het ontwerp op het masker is 0.4 (cirkels), 0.6 (vierkantjes) en 0.8 (driehoeken). Onderaan: breedte van een geïsoleerde lijn als functie van de belichtingsdosis voor verschillende ontwerpbreedtes.

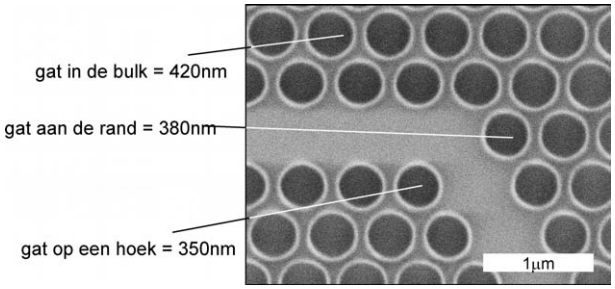

Figuur 30: Voorbeeld van optische proximiteitseffecten. De gaten aan de rand van het kristal worden kleiner afgebeeld dan de gaten in het rooster. De roosterperiode bij dit voorbeeld is $530nm$.

voldoende ruim venster om onze structuren binnen de foutenmarge van 5% te fabriceren.

5.2.2 Verschil tussen Lijnen en Gaten

Als we nogmaals figuur 29 bekijken, dan zien we dat de optimale belichtingsdosis, d.i. de dosis om een structuur op zijn ontwerpgrootte te maken, verschilt bij gaten en lijntjes. Als we beide structuren samen willen fabriceren, moeten we bij één van beiden een correctie aanbrengen. Dit is het makkelijkst bij alleenstaande lijntjes. Deze correctie moet op voorhand gekend zijn en ingecalculeerd in het maskerontwerp.

5.2.3 Optische Proximiteitseffecten

Fotonische kristallen bestaan uit dichtgestapelde gaatjes. Vermits het patroon van het masker steeds in zekere mate gediffracteerd wordt, zal het beeld van elk gat interfereren met dat van zijn buren. Dit kan zowel constructief als destructief, waardoor de gaatjes groter dan wel kleiner zullen worden. In de bulk van het rooster is dit geen probleem, omdat de procesparameters zo kunnen gekozen worden dat de gaten in de bulk juist worden afgebeeld. Aan de rand van het rooster, of bij defecten zoals golfgeleiders, hebben de gaatjes echter minder buren en kan het zijn dat ze door deze *optische proximiteitseffecten* (OPE) groter of kleiner worden afgebeeld dan in de bulk van het rooster. Figuur 30 toont een voorbeeld waarbij de gaten aan de rand merkelijk kleiner zijn dan de gaten in de bulk.

Om OPE te corrigeren moeten aanpassingen aangebracht worden op het masker. Daartoe moet men de effecten en de nodige correcties

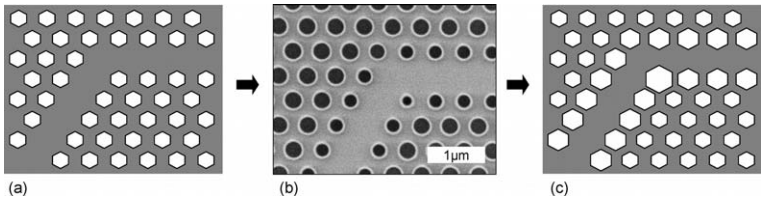

Figuur 31: Principe van optische proximiteitscorrecties voor fotonische kristalstructuren. Indien niet gecorrigeerd, zal het ontwerp op het masker (a) niet correct worden afgebeeld door optische proximiteitseffecten (b). Om dit te corrigeren, moet het masker op voorhand aangepast worden (c).

op voorhand kunnen voorspellen. Figuur 31 toont hoe de gaten op het masker groter worden gemaakt om de OPE te corrigeren.

Om dergelijke *optische proximiteitscorrecties* (OPC) te bepalen hebben we een groot aantal teststructuren aangebracht op een masker. Deze teststructuren bevatten variaties op fotonische kristalstructuren waarbij de gaatjes op de rand en de hoeken met verschillende correcties zijn aangebracht. De OPC voor nieuwe structuren kunnen dan experimenteel bepaald worden door de nodige correcties voor de best passende teststructuur rechtstreeks op te meten.

5.2.4 Verhogen van de Resolutie

Alhoewel het basis-belichtingsproces voldoende goed blijkt te presteren voor de meeste nanofotonische componenten, is het altijd interessant om over een fijnere resolutie te beschikken. Dit vergroot immers het beschikbare procesvenster (het bereik van belichtingscondities waarbinnen een structuur correct wordt afgebeeld) en vermindert ongewenste effecten als OPE.

Er bestaan verschillende technieken om de resolutie te verhogen:

- De meest voor de hand liggende techniek bestaat erin een kortere belichtingsgolflengte te gebruiken. De resolutie is immers recht evenredig met de gebruikte golflengte [50].
- Ook door de numerieke apertuur van de belichter (d.i. de maximale hoek waaronder licht nog door het lenssysteem kan propageren) te vergroten kunnen kleinere structuren worden afgebeeld. Dit gaat echter ten koste van de scherptediepte.
- Men kan betere fotolak gebruiken, die in een dunnere laag kan worden aangebracht en een betere lichtgevoeligheid heeft.

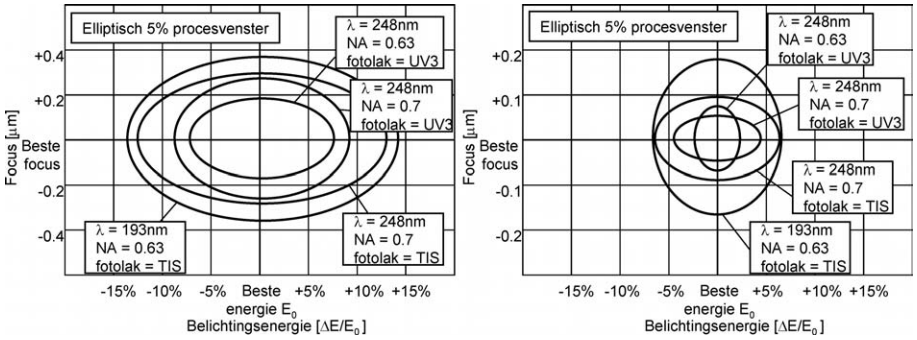

Figuur 32: Procesvenster voor vier lithografieprocessen met verschillende fotolak, numerieke apertuur en belichtingsgolflengte. Links: $300nm$ gaatjes met een periode van $500nm$. Rechts: $200nm$ gaatjes met een periode van $400nm$.

- Men kan ook de vorm en de afmetingen van de lichtbron veranderen. Dit wordt in meer detail besproken in de Engelstalige tekst.

- Geavanceerde maskers, die naast de intensiteit ook de fase van het licht beïnvloeden, kunnen de scherpte van de afbeeldingen dramatisch verhogen. Dit wordt ook in meer detail behandeld in de Engelstalige tekst.

We hebben een aantal experimenten uitgevoerd met verschillende belichtingsgolflengte ($248nm$ en $193nm$), grotere numerieke apertuur (NA) en geavanceerde fotolakken. Als maat voor de kwaliteit van het proces maten we het procesvenster als functie van de focus van de belichter en de gebruikte belichtingsdosis en dit voor een rooster van $300nm$-gaatjes met een periode van $500nm$ en voor een meer kritisch rooster van $200nm$-gaatjes met een periode van $400nm$. We bepaalden dan de grootst mogelijke ellips in het focus-energie-vlak waarvoor de gatdiameter niet meer dan 5% afweek van de ontwerpdiameter.

Dit is uitgezet in figuur 32. We zien dat er voor zowel de grote als de kleine gaatjes een belangrijke verbetering mogelijk is, zelfs zonder te moeten overschakelen naar de duurdere lithografie bij $193nm$.

5.2.5 Besluit: Diep-UV-Lithografie

We hebben aangetoond dat diep-UV-lithografie met een belichtingsgolflengte van $248nm$ de mogelijkheid biedt om nanofotonische golfgeleiderstructuren te fabriceren. Hierbij is wel grondige karakterisatie

nodig van het belichtingsproces om de nodige correcties te kunnen aanbrengen op nieuwe ontwerpen.

5.3 Etsen

Na de lithografie moeten de patronen in de fotolak overgebracht worden naar het halfgeleidermateriaal. We hebben gezien dat het vanuit optisch standpunt voordelig is om de gemiddelde brekingsindex van de oxidemantel zo laag mogelijk te houden. Daarom hebben we geprobeerd om de patronen door het oxide te etsen.

Voor de eerste experimenten maakten we gebruik van standaard SOI-plakken met een oxidedikte van $400nm$, waarbij we makkelijk door de siliciumtoplaag en het oxide konden etsen. Bij zowel de proefstructuren als de eerste functionele componenten was de etskwaliteit zeer hoog, maar het te dunne oxide veroorzaakte buitensporige substraatlek.

5.3.1 Diepe Ets

Om de hoge substraatlek te vermijden, schakelden we over op speciaal vervaardigde SOI-plakken met een bufferdikte van $1\mu m$. Het bleek echter allesbehalve makkelijk om een goed etsproces te ontwikkelen voor dergelijke dikke lagen.

Het etsen gebeurt in twee verschillende toestellen. De siliciumtoplaag wordt eerst geëtst met behulp van een $Cl_2/HBr/He/O_2$-mengsel. Voor het oxide gebruiken we daarna een etsproces op basis van een CF_4/CHF_3-mengsel. Tussen de twee etsstappen wordt de SOI-plak niet aan de lucht blootgesteld.

Het bleek mogelijk om tot $800nm$ diep in het oxide te etsen, wat vanuit optisch standpunt aanvaardbaar is. Een groter probleem is echter de ruwheden op de zijwanden die daarbij ontstaan. Die zijn zeer duidelijk zichtbaar bij de fotonische draad in de linkerhelft van figuur 33. Dergelijk grote ruwheden veroorzaken hoge propagatieverliezen, zoals we verder zullen zien.

5.3.2 Verminderen van de Ruwheden

We onderzochten verschillende technieken om de ruwheden ten gevolge van het diep etsen te verzachten. Een eerste was het behandelen van de fotolak met een plasma om de onregelmatigheden die zich in

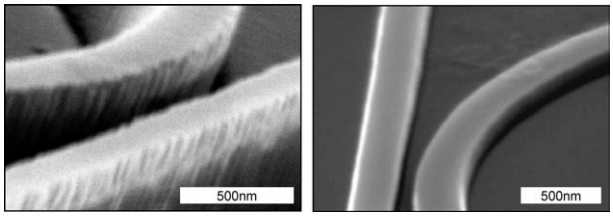

Figuur 33: Vergelijking tussen diep en ondiep etsen. De SEM-afbeeldingen tonen een detail van een fotonische draad met ringresonator. Links: diep geëtst, door zowel het silicium en het oxide. De ruwheden op de zijwand zijn duidelijk zichtbaar. Rechts: een ets van enkel de siliciumlaag. Dit veroorzaakt veel minder onregelmatigheden op de zijwand.

de fotolak bevinden, al af te vlakken. Dit gaat echter ten koste van de dikte van de fotolak, waardoor we niet meer zo diep kunnen etsen.

Een tweede mogelijkheid die we onderzochten was het oxideren van de bovenste siliciumlaag op hoge temperatuur. Omdat deze thermische oxidatie een diffuus proces is, worden de uitstekende ruwheden sneller geoxideerd en dus afgevlakt. Deze techniek heeft wel als nadeel, dat een deel van de golfgeleiderkern vervangen wordt door oxide met een lagere brekingsindex.

De oppervlakteruwheden waren echter van die aard dat geen van beide technieken in staat was ze voldoende af te vlakken.

5.3.3 Enkel de Siliciumlaag Etsen

In plaats van diep in het oxide te etsen, probeerden we dan maar om enkel de siliciumtoplaag te etsen. Door de tweede etsstap achterwege te laten, konden we de ruwheden dramatisch verminderen, zoals duidelijk wordt in figuur 33, waarbij we dezelfde ringresonator met diepe en ondiepe ets vergelijken.

5.3.4 Plasmahandeling van de Fotolak

Bij het ondiep etsen bleek echter dat de geëtste gaten steeds zowat $50nm$ kleiner waren dan de patronen in de fotolak. De lijntjes bleken breder uit te vallen. Het verschil was te groot om te compenseren met een hogere belichtingsdosis. Daarom kozen we ervoor de fotolak met een plasma te behandelen voor het etsen, waardoor een deel van de fotolak wordt aangetast en de gaatjes dus groter en de lijntjes smaller

worden. Zoals al eerder vermeld, heeft deze behandeling het voordeel dat oneffenheden in de fotolak al op voorhand afgevlakt worden.

5.3.5 Besluit: Etsen

Alhoewel het oorspronkelijk de bedoeling was om diepgeëtste structuren te maken, bleek het zeer moeilijk om de ruwheden op de zijwanden aanvaardbaar te houden. Daarom kozen we uiteindelijk voor een ets van enkel de siliciumtoplaag, waarbij we op voorhand de fotolak met een plasma behandelden .

5.4 Het Uiteindelijke Fabricageproces

Het uiteindelijke fabricageproces is geïllustreerd in figuur 34 [51]. Voor de lithografie gebruiken we een standaard proces bij $248nm$. De fotolak wordt aangebracht met een antireflectieve film en uitgebakken. Na de belichting volgt opnieuw een uitbakstap en ontwikkeling. Daarna behandelen we de ontwikkelde fotolak met een plasma. We etsen enkel de siliciumtoplaag, waarna de resterende fotolak verwijderd wordt.

5.5 Nabewerking

Om de gefabriceerde structuren uit te meten, moeten we licht inkoppelen langs gekliefde facetten. Het klieven van de grote siliciumplakken is echter geen sinecure. Daarom verdunnen we het $725\mu m$ dikke substraat tot ongeveer $250\mu m$. Daartoe delen we eerst de plak op in kleine stukken. De structuren op elk stuk worden beschermd met een fotolak en dan met was omgekeerd op een glazen plaat geplakt. Daarna wordt een deel van het substraat afgeschuurd met een suspensie van alumina. Het verdunde substraat wordt dan van de glazen plaat gehaald en gereinigd, waarna we vlakke facetten kunnen klieven.

Dit is echter een tijdrovend proces. De verwerking van een volledige plak kan makkelijk een paar dagen in beslag nemen.

5.6 Vezelkoppelaars

In deel 3 stelden we vezelkoppelaars voor op basis van een 1D-rooster. Die maken het mogelijk om licht in te koppelen vanuit een vezel die verticaal boven de structuur geplaatst is. Met dergelijke koppelaars zou de uitgebreide nabewerking voor gekliefde facetten niet meer nodig zijn.

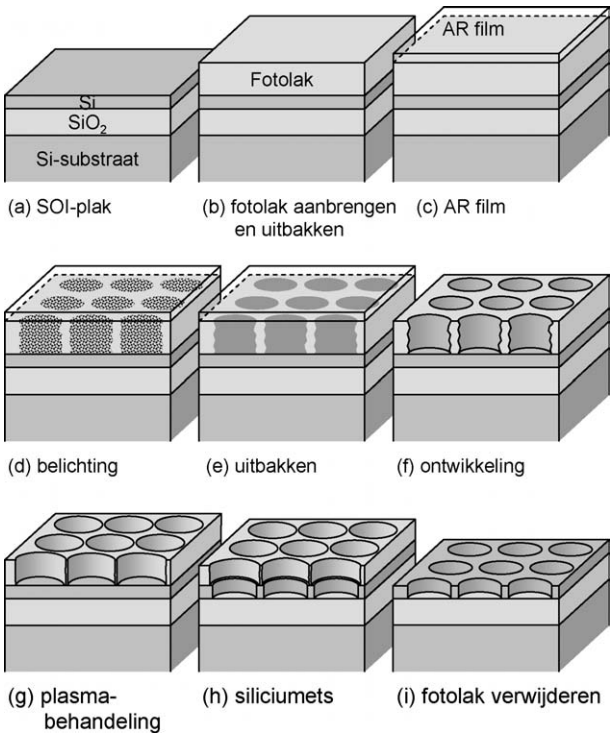

Figuur 34: Het geöptimaliseerde fabricageproces voor nanofotonische structuren in SOI.

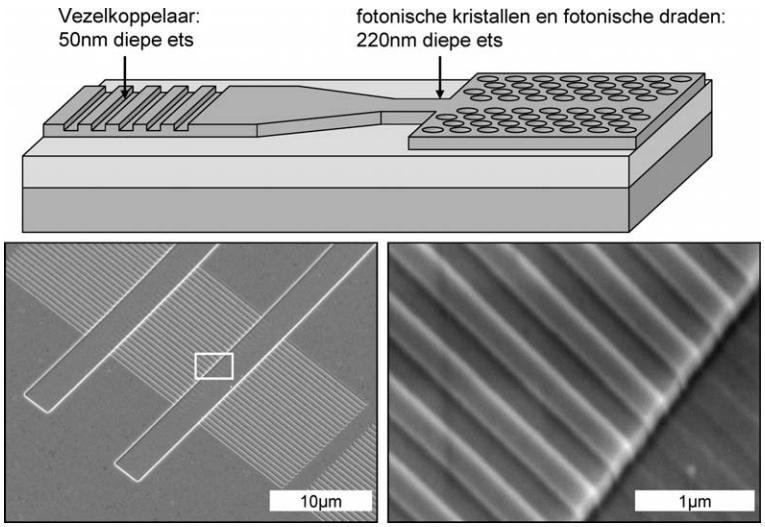

Figuur 35: Tweestapsproces. De vezelkoppelaars vereisen een andere etsdiepte dan de nanofotonische golfgeleiderstructuren.

Het probleem met de vezelkoppelaars is dat ze moeilijk samen te fabriceren zijn met nanofotonische golfgeleiders. De optimale etsdiepte voor de roosters is immers tussen $50nm$ en $80nm$, terwijl golfgeleiders minstens volledig door de siliciumtoplaag moeten geëtst worden. Dit is geïllustreerd in figuur 35.

De vezelkoppelaars moeten dus in een verschillende processtap gemaakt worden. Dit is echter geen fundamenteel probleem. De alignatie van de vezelkoppelaar met de golfgeleider is immers niet zo kritisch als die van de meeste nanofotonische componenten. Dit valt ruim binnen de mogelijkheden van de belichter. Figuur 35 toont ook een gefabriceerde structuur.

5.7 Besluit

We hebben aangetoond dat het mogelijk is om een variëteit aan nanofotonische componenten te maken met behulp van CMOS-technologie. Voor het uiteindelijke proces kozen we diep-UV lithografie bij $248nm$ en een ets van enkel de siliciumtoplaag. Dit gaf een minimum aan ruwheden op de zijwand. Om de tijdrovende nabewerking te vermijden, kunnen we vezelkoppelaars toevoegen in een tweede lithografie- en etsstap.

Het is echter wel gebleken dat een grondige karakterisatie van het fabricageproces nodig is en dit vooral bij de lithografie. Het verschil tussen lijnen en gaatjes en de invloed van optische proximiteitseffecten maken dat duidelijk.

6. Meetresultaten

Voor het karakteriseren van onze gefabriceerde structuren meten we de transmissie van het licht door de structuur. In dit deel bespreken we de meettechnieken en een aantal opmerkelijke resultaten.

6.1 De Meetopstelling

Karakterisatie van de golfgeleidercomponenten gebeurt door het inkoppelen van licht in de golfgeleider en het meten van de uitgangsrespons. Bij INTEC beschikken we daarvoor over een uitgebreide opstelling met een afstembare laser, aangepaste vezels waarmee nauwkeurig op een inkoppelfacet kan gealigneerd worden, translatietafels, vermogenmeters en een camera gevoelig voor infrarood licht.

We gebruikten twee technieken voor het in- en uitkoppelen van het licht. Enerzijds koppelen we het licht direct via een gekliefd facet in de golfgeleider. Als alternatief hebben we de vezelkoppelaars die we al eerder hebben besproken.

6.1.1 Metingen met Rechtstreekse Inkoppeling

Het licht wordt met behulp van een vezel met kleine kern en lensvormig facet ingekoppeld in de golfgeleider . Het licht in transmissie wordt opgevangen door een objectief met grote numerieke apertuur en gefocust op een vermogenmeter (figuur 36). Zo kunnen we de transmissie als functie van de golflengte opmeten tussen $1500nm$ en $1640nm$; dit is een eerder smal golflengtegebied. Om de bandstructuur van fotonische kristallen in beeld te brengen is een breedbandiger lichtbron nodig.

Inkoppeling is geen evidente zaak bij SOI golfgeleiders. De kern van de golfgeleider is zeer dun en alignatie is dus verre van evident. Daartoe zijn de vezel, het objectief en de te meten structuur op individuele translatietafels geplaatst.

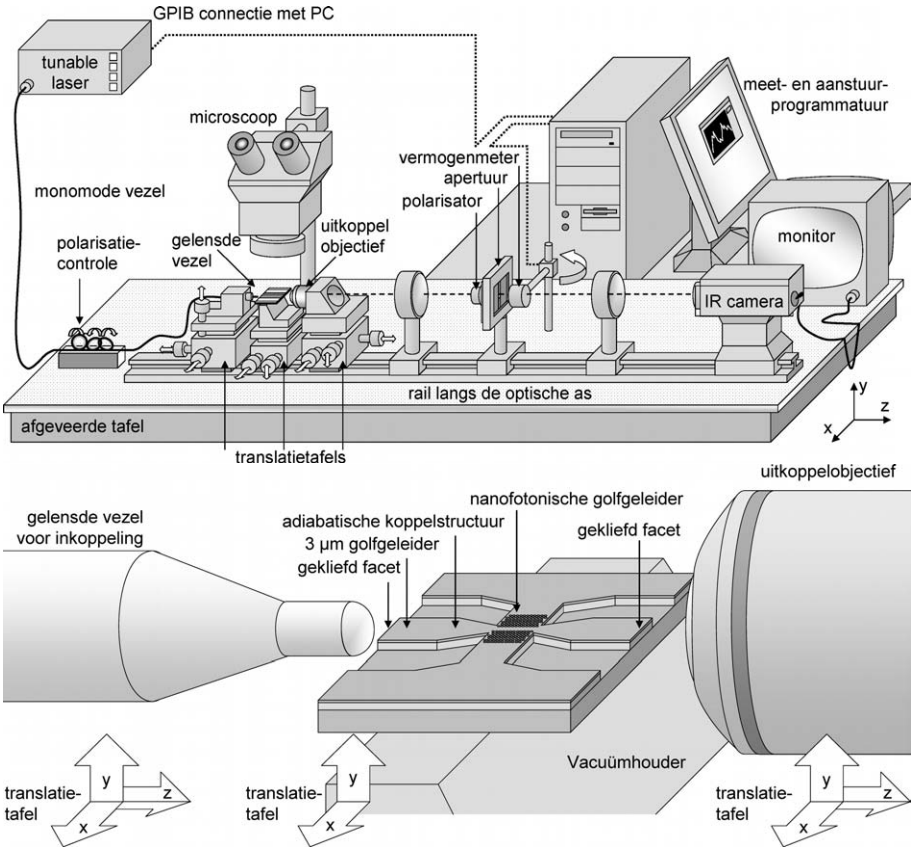

Figuur 36: De meetopstelling met directe inkoppeling. Licht van een afstembare laser wordt via een gelensde vezel in een brede golfgeleider gekoppeld. De polarisatie in de vezel wordt gecontroleerd door een torsiekracht op de vezel uit te oefenen. Het licht wordt aan de andere kant van de structuur door een objectief afgebeeld op een vermogenmeter. Deze kan eventueel weggeklapt worden zodat met een infraroodgevoelige camera het uitkoppelfacet gevisualiseerd kan worden. De laser en de vermogenmeter worden rechtstreeks door de computer aangestuurd via een GPIB-connectie. Voor nauwkeurige alignatie zijn de inkoppelvezel, de te meten structuur en het uitkoppelobjectief op individuele translatietafels gemonteerd.

6.1.2 Metingen met een Vezelkoppelaar

De verticale vezelkoppelaar is veel toleranter voor alignatie. Bovendien moet er geen facet gekliefd worden, wat een tijdrovende nabewerking van de structuren bespaart. De vezels worden eerst visueel gealigneerd met behulp van een microscoop en een camera en daarna wordt het vermogen aan de uitgang gemaximaliseerd.

Het belangrijkste nadeel van deze meettechtniek ten opzichte van directe inkoppeling is dat de vezelkoppelaars maar een bandbreedte van $60nm$ hebben waarbinnen efficiënt wordt ingekoppeld. Voor metingen over een groter bereik gebruiken we directe inkoppeling.

6.2 Meten van Propagatieverlies

Voor het meten van propagatieverliezen hebben we een aantal meettechnieken tot onzer beschikking. Deze berusten steeds op de zogenaamde *cut-back*-techniek. Daarbij wordt het transmissieverlies gemeten van golfgeleiders van verschillende lengte. Als de verliezen uitgezet worden op logaritmische schaal (in dB) als functie van de lengte van de golfgeleider(in mm), liggen ze normaliter op een rechte. De helling van die rechte geeft ons dan de propagatieverliezen in dB/mm.

Het bepalen van het verlies van een enkele golfgeleider kan op verschillende wijzen gedaan worden, die uitgebreid besproken worden in de Engelstalige tekst. De meest eenvoudige is rechtstreeks het vermogen aan de uitgang meten. Dit uitgangsvermogen is echter afhankelijk van een aantal externe factoren, zoals de nauwkeurige alignatie van de inkoppelvezel. Bij directe inkoppeling krijgen we bovendien ook oscillaties als functie van de meetgolflengte omdat de gekliefde facetten een Fabry-Perotcaviteit vormen.

We kunnen de oscillaties van een Fabry-Perotcaviteit echter gebruiken om het verlies binnen de caviteit te bepalen. De transmissie van een Fabry-Perotcaviteit is immers een quasi-periodieke functie van de golflengte, waarbij de periode afhangt van de optische lengte van de caviteit. Het contrast tussen de maxima en de minima is een maat voor het verlies. Hieruit kunnen we veel nauwkeuriger de transmissie bepalen dan rechtstreeks uit het opgemeten vermogen.

Deze techniek is echter enkel bruikbaar als er binnen de structuur geen extra reflecties optreden. Bij de overgang tussen een fotonische draad en een fotonisch kristal is er echter steeds wat reflectie. Daardoor krijgen we een samenstelling van verschillende Fabry-Perotcaviteiten. Het transmissiespectrum wordt dan een superpositie van periodieke

Generatie		SOI		Ets		Maskers				Opmerkingen
		Si	Ox	Si	Ox	1	2	3	4	
1	maa/02	205	400	205	400	X				Dun ox.
2	feb/03	220	1000	220	800	X	X			Ruwheden
3	jun/03	220	1000	220	geen		X			
4	aug/03	220	1000	220	geen			X		Tweestaps
5	dec/03	220	1000	220	geen				X	Tweestaps

Tabel 1: Overzicht van gefabriceerde nanofotonische golfgeleiders in SOI, met de basisparameters van het fabricageproces. Laagdiktes en etsdieptes zijn uitgedrukt in nm.

structuren. Door het fouriertransformeren van het transmissiespectrum kunnen we de periodieke respons van de verschillende caviteiten echter onderscheiden, vermits de periode van elke caviteitsrespons omgekeerd evenredig is met de caviteitslengte. Uit de relatieve hoogte van de verschillende pieken in de fouriergetransformeerde van het transmissiespectrum, kunnen we de caviteitsverliezen extraheren. Als we deze techniek lokaal toepassen op een klein venster van het volledige meetspectrum, kunnen we bovendien de verliezen als functie van golflengte berekenen.

6.3 De Gefabriceerde Componenten

Zoals eerder besproken, hebben we verschillende generaties nanofotonische componenten gefabriceerd. Deze werden gegroepeerd op de maskers die nodig waren voor de diep-UV-lithografie. Deze maskers worden besproken in appendix A. De verschillende generaties zijn opgesomd in tabel 1 met hun specifieke fabricageparameters.

6.4 Generatie 1: SOI met een Dunne Buffer

Na succesvolle lithografie- en etsproeven hebben we een aantal structuren gemaakt op basis van ontwerpen die beschikbaar waren in de literatuur [7, 52]. Door de dunne buffer hadden deze golfgeleiders echter zeer hoge propagatieverliezen ten gevolge van substraatlek. Voor een fotonische draad van $500nm$ breed maten we $130dB/mm$ en minder naarmate de golfgeleider breder was. Zoals we gezien hebben, is dit omdat bij smallere golfgeleiders de mode zich verder in de buffer uitstrekt. Bij een W1 fotonische kristalgolfgeleider maten we $100dB/mm$ waar de mode zich onder de lichtlijn bevindt.

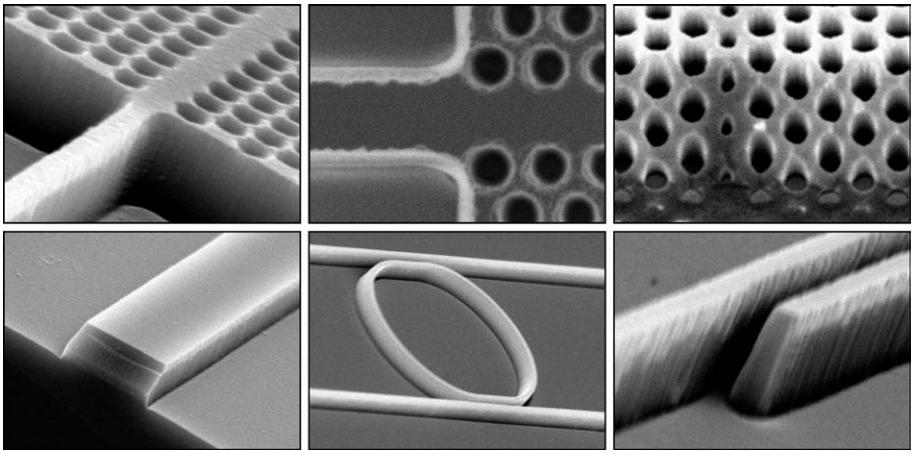

Figuur 37: Nanofotonische structuren in SOI met een dikke oxide buffer en een diepe etsstap. Bovenaan: een fotonische draad gekoppeld aan een W1 fotonische kristalgolfgeleider, hetzelfde in bovenaanzicht, een schuin gekliefde fotonische kristalgolfgeleider met kleine defectgaatjes. Onderaan: een gekliefd facet van een $3\mu m$ brede golfgeleider, het begin van en directionele koppelaar, een 'racetrack' resonator. Bemerk de duidelijk zichtbare ruwheden op de zijwanden.

6.5 Generatie 2: Diepgeëtste Golfgeleiders

Voor de nieuwe SOI-plakken met een dikke buffer ontwikkelden we een diep etsproces. Zoals echter duidelijk blijkt uit figuur 37, zijn de zijwanden allesbehalve glad. Alhoewel de substraatlek nu verwaarloosbaar zou moeten zijn, verwachten we toch aanzienlijke verliezen ten gevolge van verstrooiing aan deze ruwheden.

Figuur 38 toont de propagatieverliezen van fotonische draden met verschillende breedtes. De individuele punten geven de interne caviteitsverliezen die werden bepaald door het transmissiespectrum te fouriertransformeren en de hoogte van de pieken te meten. De helling geeft een propagatieverlies van $34dB/mm$ voor een fotonische draad van $500nm$ breed. Opnieuw nemen de verliezen af voor smallere draden, ditmaal omdat de veldsterkte op de wand kleiner is.

Voor fotonische kristalgolfgeleiders wilden we zowel de modi met een even symmetrie als met een oneven symmetrie opmeten. Daarom hebben we de golfgeleiders geïmplementeerd in twee varianten: één met een gewone inkoppelgolfgeleider om de even mode aan te slaan en een inkoppelstructuur die de even inkoppelmode omzet naar een mengeling van een even en een oneven mode.

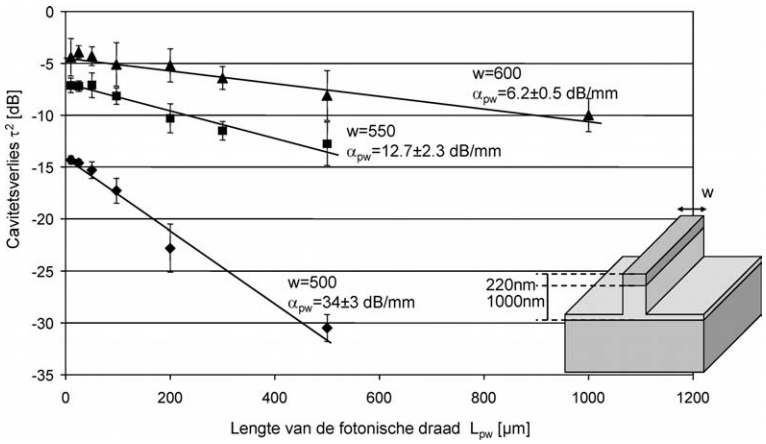

Figuur 38: Propagatieverliezen voor diepgeëtste fotonische draden met verschillende breedtes.

Figuur 39 geeft de propagatieverliezen en de resterende caviteitsverliezen van een fotonische kristal, berekend uit de piekhoogte van de lokale fouriergetransformeerde van het opgemeten spectrum. We vinden een minimale transmissie van $21dB/mm$ rond een golflengte van $1515nm$. Vermits dit een meting betreft met gewone inkoppeling, gaat het hier om een mode met even symmetrie.

Alhoewel deze verliezen beduidend lager zijn dan voor een dunne oxidebuffer, zorgt de verstrooiing aan de ruwe zijwanden nog voor hoge propagatieverliezen.

6.6 Generatie 3: Enkel de Toplaag Geëtst

Vermits het moeilijk bleek om de oxidebuffer te etsen zonder ruwheden te creëren, hebben we ervoor gekozen enkel de toplaag te etsen. Figuur 40 toont een aantal structuren die met dit proces vervaardigd zijn. We zien dat de zijwanden beduidend gladder zijn dan bij de diepgeëtste structuren.

Figuur 41 toont de propagatieverliezen van fotonische draden voor verschillende breedtes. Voor deze structuren konden we de verliezen makkelijk extraheren uit het contrast van de franjes van het Fabry-Perot-transmissiespectrum. We zien dat voor de $500nm$-brede golfgeleiders de verliezen bijzonder laag zijn: $0.24dB/mm$. Voor smallere golfgeleiders nemen de verliezen wel sterk toe.

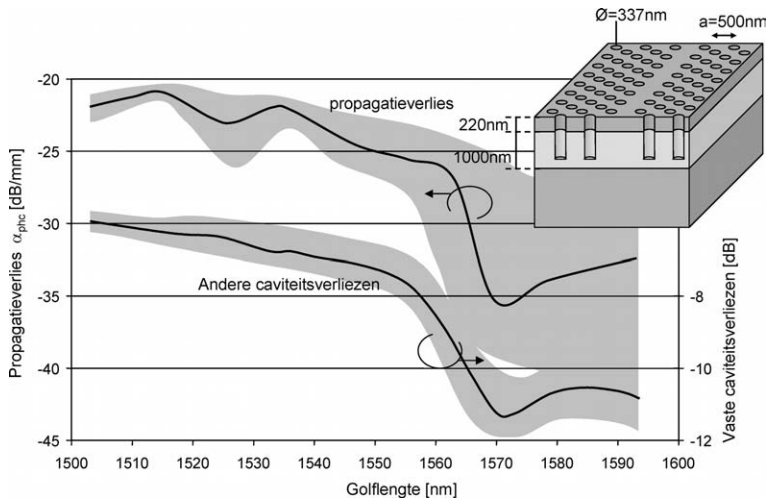

Figuur 39: Propagatieverliezen van een diepgeëtste W1 fotonische kristalgolfgeleider met een roosterperiode van $500nm$ en een gatdiameter van $337nm$. De vaste caviteitsverliezen zijn eveneens aangeduid.

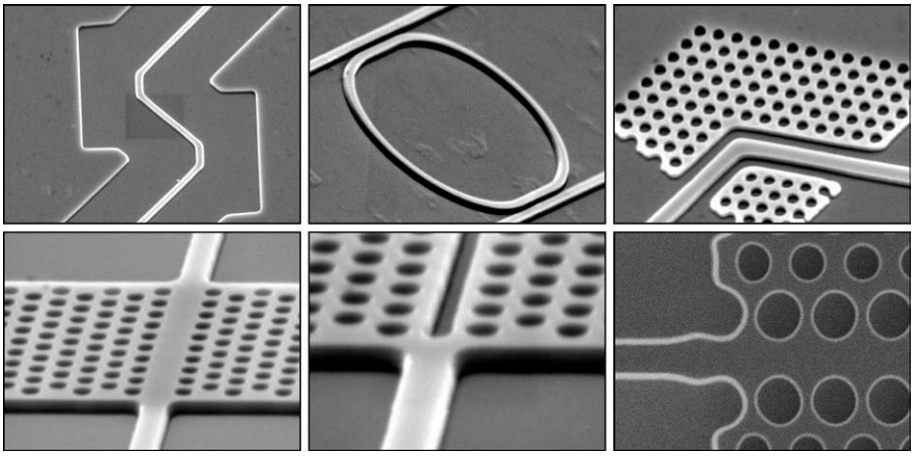

Figuur 40: Nanofotonische structuren in SOI met enkel een siliciumets. Bovenaan: een fotonische draad met 2 hoekspiegels, een 'racetrack' resonator en een draad ingebed in een fotonisch kristal. Onderaan: een W1 fotonische kristalgolfgeleider, een W1 golfgeleider met een enkele sleuf, een W1 golfgeleider met grotere gaten op de rand.

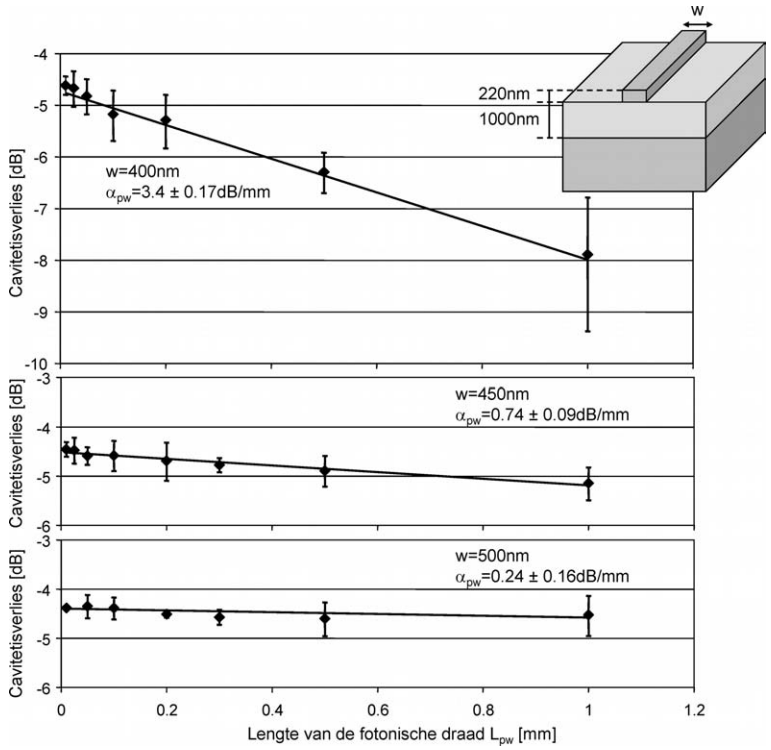

Figuur 41: Propagatieverliezen voor fotonische draden met enkel een siliciumets.

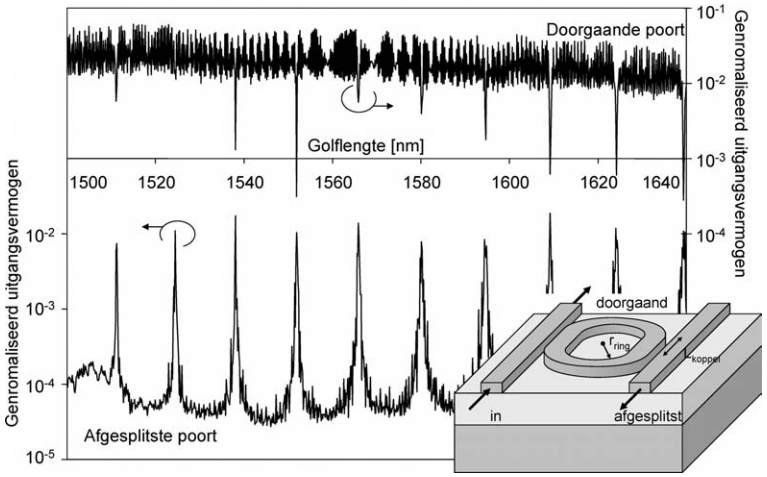

Figuur 42: Transmissie van een 'racetrack' resonator in de doorgaande poort en de afsplitste poort. De resonator heeft een kwaliteitsfactor Q van meer dan 3000 en een koppelefficiëntie van 80%.

Door deze lage propagatieverliezen kunnen we ook goede resonatoren maken met deze fotonische draden. Een meetresultaat van de resonator met 'racetrack' configuratie uit figuur 40 is weergegeven in figuur 42. We kunnen duidelijk zien waar er resonantie optreedt en het licht van de ingang wordt afgesplitst langs de resonator. De kwaliteitsfactor Q is groter dan 3000 en de overkoppelefficiëntie is 80%. We hebben eveneens ringresonatoren uitgemeten met een Q van 8000, weliswaar met een lagere overkoppelefficiëntie.

Ook de fotonische kristalgolfgeleiders hebben lagere verliezen. De propagatieverliezen van een W1 golfgeleider met een asymmetrische inkoppeling zijn weergegeven in figuur 43. Rond $1525nm$ bedragen de verliezen van de oneven mode slechts $7.5dB/mm$.

De structuren waarbij enkel het silicium geëtst is presteren duidelijk beter dan de diep geëtste structuren. De fotonische draden hebben uiterst lage propagatieverliezen en ook voor de fotonische kristalgolfgeleiders zijn de verliezen al merkelijk lager.

6.7 Generatie 4: Koppelstructuren

Voor de eerste structuren met een siliciumets hebben we de verliezen opgemeten met directe inkoppeling. We hebben op deze golfgeleiders

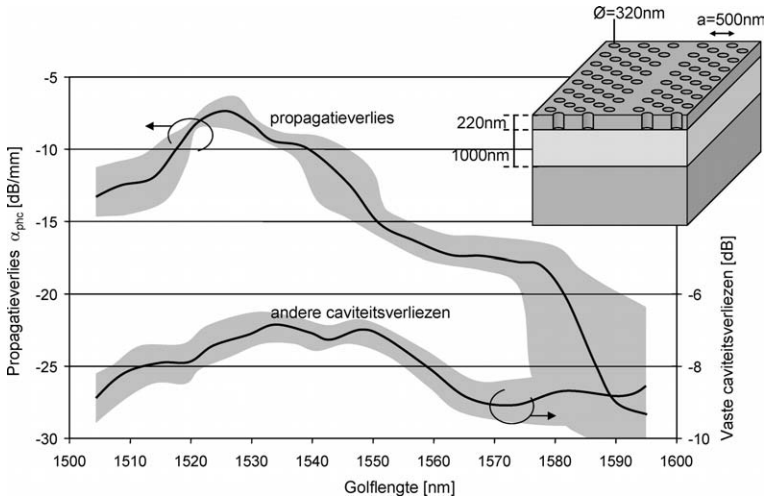

Figuur 43: Propagatieverliezen van een W1 fotonische kristalgolfgeleider met enkel een siliciumets. Het rooster heeft een periode van $500nm$ en de gaatjes meten $320nm$. De verliezen werden berekend met behulp van de piekhoogte in een lokale fouriergetransformeerde van het transmissiespectrum.

ook vezelkoppelaars aangebracht met een tweestapsproces[11]. Met hetzelfde proces hebben we ook interferentiekoppelaars vervaardigd om het licht van de brede golfgeleiders naar een fotonische draad te koppelen[12]. Een aantal van deze structuren zijn afgebeeld in figuur 44.

We maten de efficiëntie van de vezelkoppelaars door de transmissie van vezel naar vezel op te meten in een structuur met een $7mm$ lange brede golfgeleider die in het midden adiabatisch versmalt tot $500nm$ om alle modi behalve de grondmode weg te filteren. Als we deze golfgeleider verliesloos veronderstellen kunnen we een ondergrens voor de koppelefficiëntie schatten. Figuur 45 geeft de koppelefficiëntie van een enkele vezelkoppelaar als functie van golflengte. Net als in het model krijgen we een gaussiaans verloop, met een maximum van 21% en een $3dB$ bandbreedte van $60nm$.

De vezelkoppelaars hebben we ook gebruikt om koppelstructuren uit te meten. Figuur 46 toont de transmissie, van vezel naar vezel, voor de verschillende structuren uit figuur 44. We zien dat de interferentiekoppelaar met een lengte van $15.4\mu m$ gelijkaardige prestaties neerzet

[11] De ontwikkeling van deze roosterkoppelaar is het werk van ir. Dirk Taillaert, uitgevoerd in het kader van zijn doctoraatsonderzoek.

[12] De ontwikkeling van deze interferentiekoppelaar is het werk van ir. Bert Luyssaert, uitgevoerd in het kader van zijn doctoraatsonderzoek.

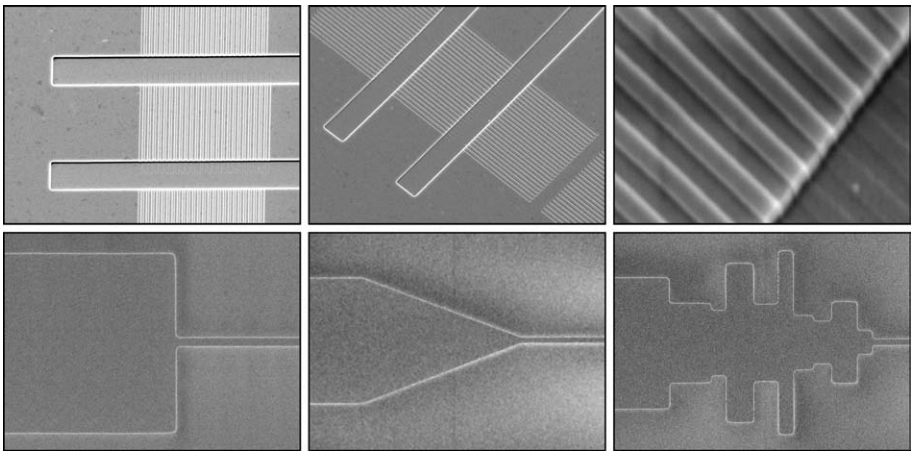

Figuur 44: Vezelkoppelaars en compacte koppelstructuren. Bovenaan: 1D vezelkoppelaar met een ondiepe ets samen met een golfgeleider die in het silicium geëtst is. Onderaan: Directe koppeling van brede naar smalle golfgeleider, een korte lineare koppelstructuur, een koppelstructuur op basis van interferentie.

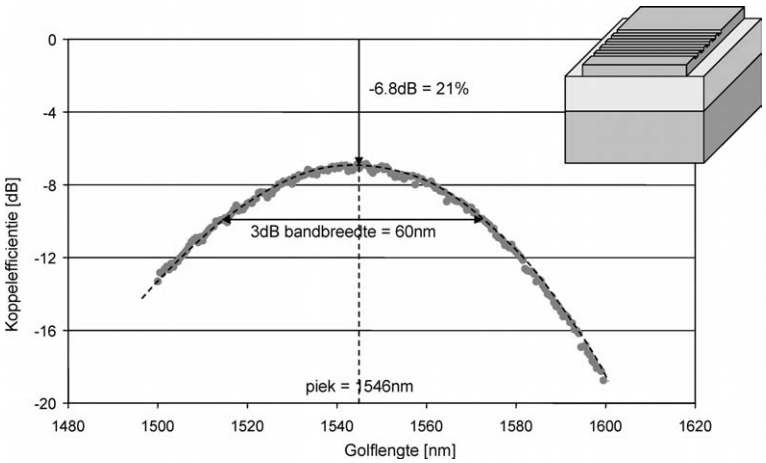

Figuur 45: Koppelefficiëntie van een 1D vezelkoppelaar. De curve is gaussiaans met een maximum koppelefficiëntie van $-6.8dB$ of 21%, en een $3dB$ bandbreedte van $60nm$.

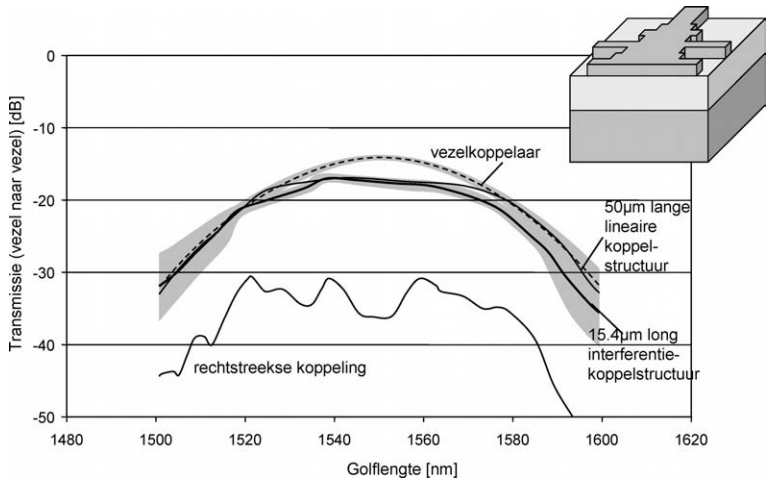

Figuur 46: Transmissie van compacte koppelstructuren. De grafiek vergelijkt de vezel-naar-vezel transmissie van een $7mm$ lange golfgeleiderstructuur bestaande uit een $10\mu m$ brede golfgeleider met een koppelstructuur naar een $560nm$ brede golfgeleider en terug. Inkoppeling en uitkoppeling gebeurt met behulp van verticale vezelkoppelaars. De referentiecurve is die van de vezelkoppelaar met adiabatische koppelstructuren. De opgemeten structuren staan afgebeeld in figuur 44: een $50\mu m$ lange lineaire koppelstructuur, directe koppeling en een geoptimaliseerd interferentiekoppelaar bestaande uit 10 secties.

als een lineaire koppelstructuur die drie maal zo lang is. De overdracht van de brede naar de smalle golfgeleider is 70% over een golflengtebereik van $50nm$.

6.8 Besluit

We hebben een groot aantal componenten uitgemeten, vervaardigd met verschillende fabricageprocessen. Zoals we verwachtten, verkregen we de beste resultaten, d.i. de laagste propagatieverliezen, met structuren waarbij enkel de siliciumtoplaag geëtst is. We vergelijken onze resultaten met verliesmetingen uit de literatuur [48] in tabel 2. In onze structuren, vervaardigd met diep-UV-lithografie, maten we propagatieverliezen van $0.24dB/mm$ voor fotonische draden en $7.5dB/mm$ voor planaire fotonische kristalgolfgeleiders.

De verticale vezelkoppelaars hadden een efficiëntie van meer dan 20% en de interferentiekoppelaars hadden 70% transmissie. Een van de belangrijkste beperkingen voor de metingen is de bandbreedte van de lichtbron. Voor golflengte-afhankelijke structuren zoals fotonische

	$500nm$ draad	W1 golfgeleider
Deep etch	34 dB/mm	21 dB/mm
Si-only etch	0.24 dB/mm	7.5 dB/mm
Literatuur*	0.16 dB/mm	2.4 dB/mm
* Vervaardigd met elektronenbundellithografie [48]		

Tabel 2: Vergelijking van de propagatieverliezen in nanofotonische golfgeleiders.

kristallen is het belangrijk om het gedrag over een groot golflengtegebied te kennen.

7. Besluit

De basiscomponenten voor nanofotonische golfgeleidercircuits zijn *fotonische draden* en *fotonische kristallen*. Alhoewel fotonische kristallen met hun PBG zeer interessante eigenschappen hebben, blijkt het toch niet eenvoudig om goede golfgeleiders te ontwerpen. Het golfgeleidingsmechanisme van fotonische draden berust daarentegen uitsluitend op het hoge brekingsindexcontrast tussen kern en mantel. Dit maakt de structuren veel eenvoudiger te begrijpen en te modelleren. Ze zijn dan ook betere kandidaten voor eenvoudige golfgeleiding en interconnectie.

De belangrijkste eis voor een goede golfgeleider is een laag propagatieverlies. We hebben de verschillende verliesmechanismes in zowel fotonische draden als planaire fotonische kristalgolfgeleiders besproken. Daaruit blijkt dat fotonische draden opnieuw de voorkeur genieten, als we het belangrijkste verliesmechanisme, nl. verstrooiing aan ruwheden, kunnen beperken. Dit vereist een goed fabricageproces.

Voor fotonische kristallen hebben we de verstrooiingsverliezen bekeken als functie van de verticale lagenstructuur. Zelfs in een perfect vervaardigd fotonisch kristal kan er immers verstrooiing optreden. Voor deze intrinsieke verstrooiing zijn er twee goede regimes. Bij een laag verticaal contrast wordt het licht over gans de structuur weinig verstrooid, maar de verliezen accumuleren wel over de ganse propagatielengte. Bij een hoog verticaal brekingsindexcontrast kunnen we een verliesloze geleide mode opwekken, maar bij een verstoring van de symmetrie van de golfgeleider kunnen de verstrooiingsverliezen wel hoog oplopen.

Voor de fabricage van onze componenten maakten we gebruik van de faciliteiten van IMEC in Leuven en in het bijzonder van diep-UV-

lithografie bij $248nm$. Silicium-op-isolator was hiervoor het geschikte basismateriaal. Het blijkt dat we inderdaad nanofotonische structuren kunnen maken met CMOS-technologie, maar dat grondige proceskarakterisatie vereist is, zeker met het oog op optische proximiteitseffecten in fotonische kristallen. Voor het etsen kozen we er uiteindelijk voor om enkel de silicium-toplaag te etsen, omdat bleek dat een diepe ets onoverkomelijke ruwheid op de zijwanden veroorzaakt.

We hebben de vervaardigde structuren gekarakteriseerd door de transmissie als functie van golflengte te meten. De grootste beperking werd hierbij gevormd door de bandbreedte van de lichtbron. Voor fotonische draden vonden we uiterst lage verliezen van $0.24dB/mm$ bij een golfgeleiderbreedte van $500nm$. Tevens hebben we ringresonatoren met een kwaliteitsfactor Q van 8000 opgemeten.

De fotonische kristalgolfgeleiders hadden significant hogere verliezen dan de fotonische draden. Voor een W1 golfgeleider maten we $7.5dB/mm$.

Naast golfgeleiders hebben we ook koppelstructuren gemaakt om licht van en naar monomode optische vezels te koppelen. Roosterkoppelaars koppelen licht van een verticale optische vezel naar een $10\mu m$ brede SOI golfgeleider met een efficiëntie van 21%. Met een interferentiekoppelaar konden we dan het licht naar een $550nm$ brede fotonische draad koppelen met een koppelefficiëntie van 70%.

Om dit werk in weinig woorden samen te vatten, kunnen we stellen dat we werkbare route hebben blootgelegd voor de commerciële toepassing van nanofotonische golfgeleiders die op grote schaal geïntegreerd kunnen worden. Fotonische draden hebben aanvaardbaar lage verliezen en fotonische kristallen bieden met hun dispersief gedrag een waaier van functionele elementen aan. Vezelkoppelaars en compacte koppelstructuren maken interactie mogelijk tussen deze componenten en de buitenwereld. Voor de fabricage baseren we ons op mature CMOS-technologie, die zeer geschikt is voor de bewerking van SOI. Zelfs de elektronica-industrie ziet optische interconnecties als een toekomstige uitbreiding voor CMOS-componenten. Nanofotonische golfgeleiders in silicium zoals we die hier hebben gedemonstreerd zijn een goede kandidaat om dit waar te maken.

English Text

Preface

This work is on *nanophotonics*. It's about the manipulation of light on a tiny scale. Light that carries information all over the planet. Light for fibre-optic communications. Five years ago, the *dot.com* boom was breaking loose, and for telecommunication, the sky was not the limit, but merely a stop on the way to higher ambitions. Nanophotonics, including photonic crystals, would make it possible to reduce current photonic devices to very small dimensions and add an uncountable number of functions onto a chip, just like the electronics industry had done two decades earlier. In the early twenty-first century everyone would be surfing the internet through its own high-speed optical fibre link. Not yet so. *Dot.com* evolved into *dot.came* and many ambitions (and money) were stored on liquid nitrogen. *More* has been replaced by *cheaper*. Luckily, nanophotonics also carries a lot of cost-reduction promises, so this business is still running, be it in a more temperate climate.

When I started my work as a PhD student, I was the first person in our group to break open the young research area of photonic crystals. While there was already a considerable experience with periodic structures, the use of high-contrast photonic crystals for waveguiding purposes was then mostly unexplored territory. At that time, photonic crystals would solve all the problems of the photonics industry. Because of their amazing properties, it was possible to create complex devices from elementary building blocks with the ease of stacking Lego. While this simple concept still stands, it proved not so straightforward to design the basic Lego block in the first place.

In the past five years, photonic crystal research has become a prominent activity in the group, with at least five people working on photonic crystal-related research topics. Over the years, we have developed a thorough understanding of the intricacies of photonic crystal optics. This is one of the purposes of this book: to share these insights.

With this in mind, we designed an alternative cover page for this book, which you find somewhere on the next pages.

As with much research work, the results in this book could not be brought about by one man working on his own for eight hours a day. Many of the ideas are a result of fruitful discussions with colleagues both from within the group and other research groups. At the risk of forgetting someone, I would like to mention a number of people by name.

Thank you, Dirk and Bert, for finding ways to get light into my waveguides, and for some great moments, especially when we were 'accidentally' stranded in San Francisco, 10 days early for the PECS conference in LA.

Thank you, Pieter, Joris and Hans. Supervising you as thesis students was a pleasure, and parts of your work and many valuable insights are incorporated in this book. An additional thanks to Pieter, for taking much measurement and simulation load off my shoulders when too many tasks were piling up.

Thank you, Peter and Lieven, for CAMFR, and all the support that was needed while I got the opportunity to thoroughly 'test' the program. Without your work, at least a whole chapter would be missing.

Thank you, Roel, for creating the opportunity to do this work, for the valuable discussions, with both enthusiasm and criticism, and of course for the suddenly appearing can-you-do-this-by-yesterday deadlines, which kept me on my toes.

Thanks to my office-mates, Bart, Freddie, Watje, Hannes, and recently Pablo. Together we have slaughtered thousands of electronic worms. Thanks also to building 39 for not being torn down in all those years, and its inhabitants (including myself) who sometimes turned the staircase into a battle-zone. As a whole, the OCS group, recently renamed to *photonics* group, formed a group of colleagues with whom it was fun not only during work, but also after hours, for a drink in the Vooruit or Backstage, or at the yearly *Opto*-day.

To expand the circle somewhat further, I would like to thank the people at INTEC who are constantly working behind the scenes in order to let us concentrate on the things that are important to us (like the aforementioned Opto-day).

This work would have been nigh impossible without the people of IMEC in Leuven, who, in order to get our photonic crystals made, processed zillions of wafers (sometimes with the wrong process, though). Thank you, Vincent, Johan, Stephan, Diziana, Rudi en Johan.

Also, the partners of the PICCO project deserve a big thank you. The three years of this European project were extremely stimulating, and also far too short. The projects meetings were, well, fun. Much of my work has been carried out in the framework of this project, and without the collaborations, things would have been very different.

Also my sincere gratitude to everyone who has found at least one typo or goof-up in this book (and preferably before it was sent to the printers).

Outside of my work, there are some other people who have kept me on my feet in those years. In some ways, they have contributed to my work, too. Paps and Mams, thanks for all the support (including financial) that helped me through university. Thank you, little Goedele, for forcing my mind off the job every evening. And most of all, thank you, Katrien. You were always there for me, still are, and I hope you will always be.

Wim Bogaerts
Gent, 12 December 2003

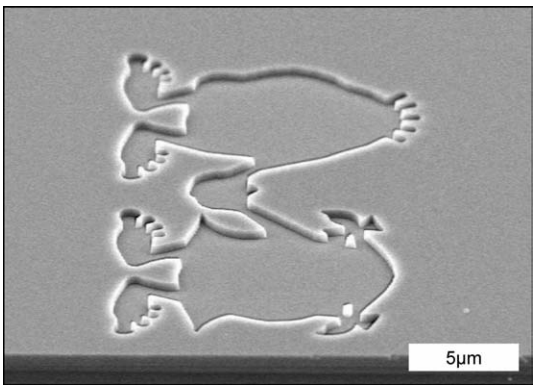

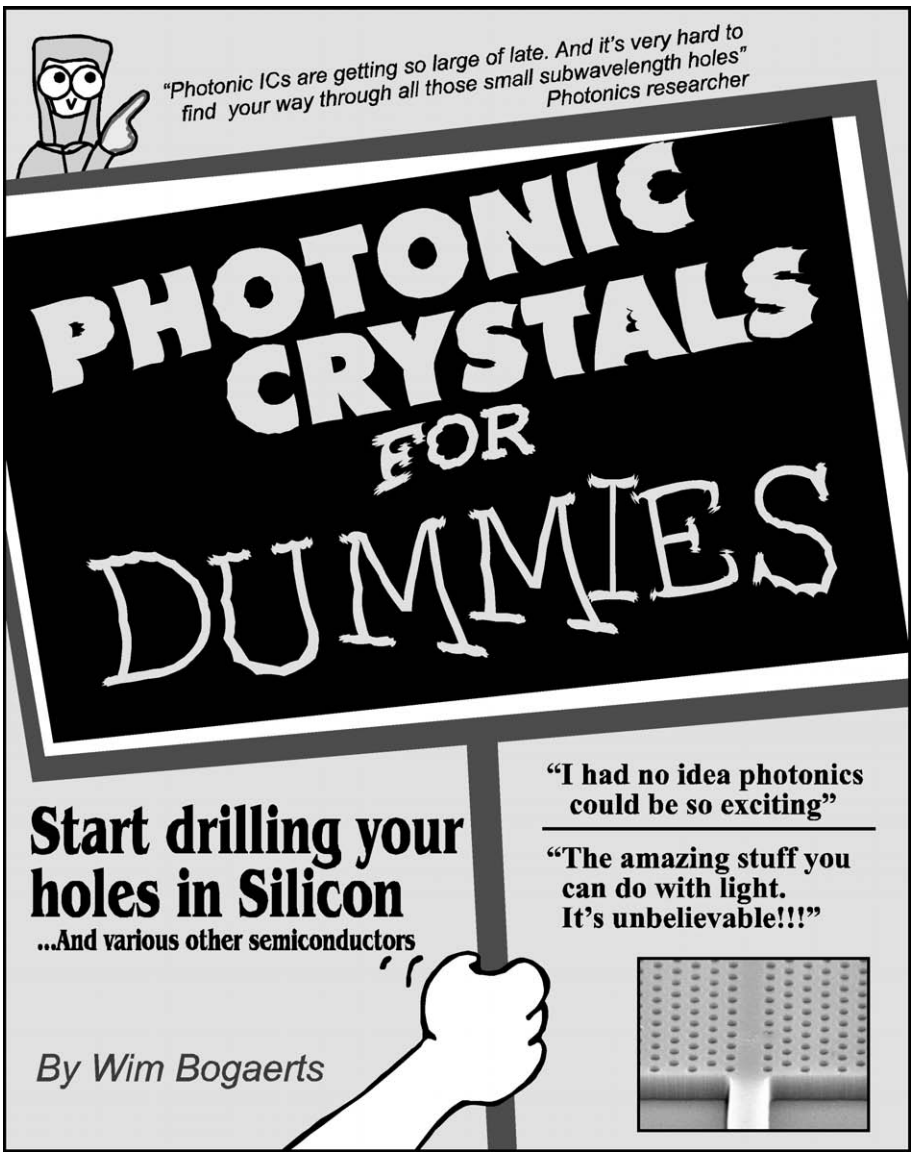
"Photonic ICs are getting so large of late. And it's very hard to find your way through all those small subwavelength holes"
Photonics researcher
PHOTONIC CRYSTALS FOR DUMMIES
Start drilling your holes in Silicon
...And various other semiconductors
By Wim Bogaerts
"I had no idea photonics could be so exciting"
"The amazing stuff you can do with light. It's unbelievable!!!"

Contents

Contents

Chapter 1

Introduction

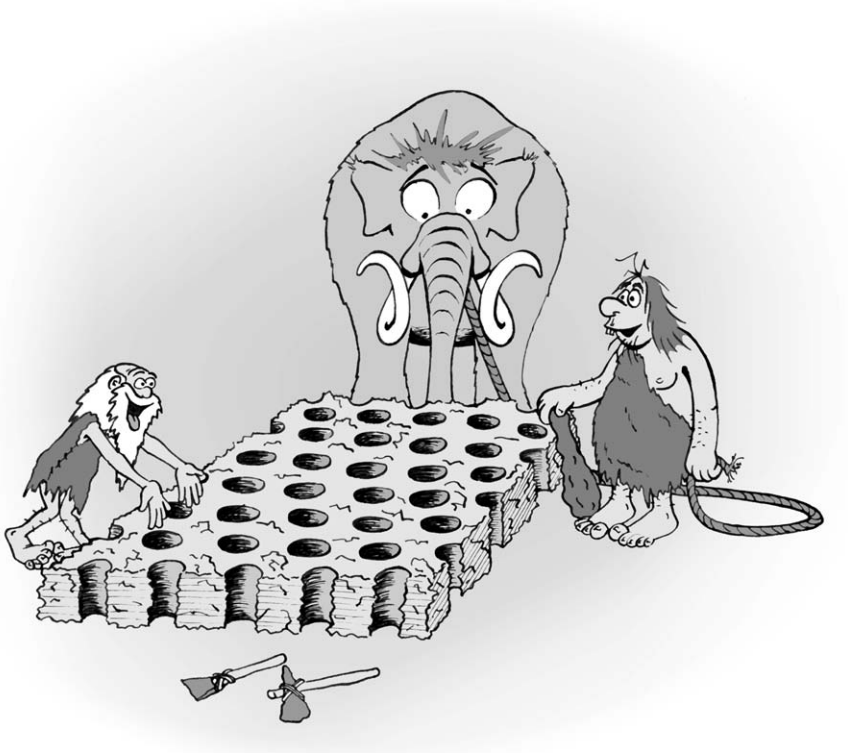

"Look, Zog! This will be the next breakthrough in telecommunications!"

Chapter 1

Introduction

In this work we will present planar nanophotonic waveguides, including photonic crystals, in Silicon-on-insulator. However, before we go through the intricate details of the theory, design and fabrication process, we will try to clarify the rationale behind this work. The drive for ultra-compact photonics is mainly generated in the telecom business. Integration and cost-reduction is one of the main drives behind the development of ultra-compact photonics, also dubbed *nanophotonics*.

In this introductory chapter, we will briefly discuss the use of photonics in telecommunication networks, as well as introduce *photonic crystals*, important candidates for advanced photonic components.

1.1 Telecom, Internet and Photonics

In the last few centuries, telecommunication has reshaped the way we think. While previously information flow was mainly restricted to the speed by which people moved from A to B, the telegraph, then telephone and more recently the internet has allowed us to contact anyone, anywhere, at any time, or retrieve vast amounts of information from a variety of sources.

The first telecommunication links consisted of electrical links. Electrical communication is still used in our daily life in the form of telephone, television, computer networks etc. However, electrical transmission doesn't scale up to the data rates needed for long-distance communication backbones.

Optical fibres, made of ultra-pure glass can guide light over large distances with little propagation loss. In one decade, fibres have superseded electrical connection throughout most of the backbone infrastructure. For local connections and distribution, conventional electrical links are still dominant, mainly because of the cost of replacing the installed base of twisted-pair and coaxial cables in each street, home and office building.

While optics are extremely suitable for signal transport, they are hard hard to use for logical circuits. In telecommunication networks, routing and regeneration is still done electronically. The resulting need for multiple *optoelectronic* conversions seriously limits the growth of fibre-optic communications.

Optical fibres seem to have an almost limitless capacity to carry information. In the infrared wavelength ranges around $1.55\mu m$ and $1.3\mu m$, optical fibres have very low absorption and/or good dispersion characteristics, which allows the transport of high bit rate signals. However, over distance the signal degrades, and this effect gets stronger with higher bit rates. Therefore, instead of using a single broadband channel, multiple narrow-band channels at different wavelengths are used. Because the wavelengths propagate independently, many channels can be used simultaneously, boosting the capacity of the fibre. This is called *wavelength division multiplexing* (WDM). This technique can be applied to many existing point-to-point fibre links already installed by replacing the transmitters and receivers at both ends. However, because all wavelength channels need to be processed individually, the processing equipment in the end points becomes increasingly complex. Not only do we need to split up the various wavelengths, but each wavelength will need its own optoelectronic and electronic processing circuitry. This multiplies the burden of having to do all functionality in the electrical domain.

One of the key objective of *photonics* is the development of optical alternatives that don't need the costly electro-optic conversions. Just like in electronics, photonic components can be integrated on-chip resulting in a *photonic integrated circuit* (PIC). In such a PIC, the routing of light is performed by optical *waveguides*. However, waveguides currently used in commercial PICs pose certain limits to the reduction of photonic ICs. New technologies, especially to replace the bulky waveguides, are needed to make further improvement possible.

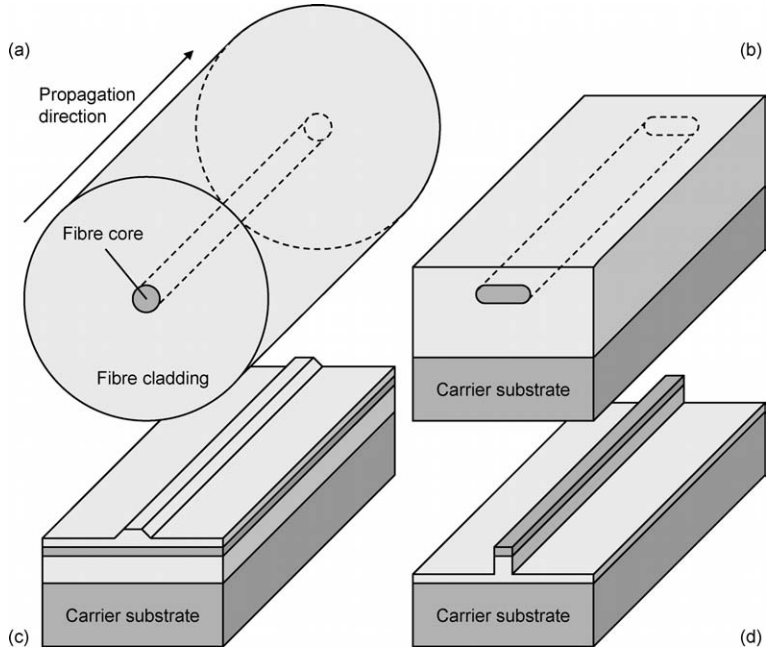

Figure 1.1: Various waveguide cross sections. (a) A single mode optical fibre, with a glass core of $9\mu m$ and a glass cladding of $125\mu m$ diameter. The index contrast between core and cladding is typically lower than 1%. (b) A Silica-on-Silicon waveguide, with similar dimensions and index contrasts. (c) A typical low-contrast III-V semiconductor ridge waveguide. The index contrast between core and cladding layers is typically in the range of 3.2 to 3.4, with a ridge width of several microns. (d) A high-contrast Silicon-on-insulator photonic wire with submicron dimensions. The core material is Silicon with an index of 3.45, compared to a cladding of SiO_2 with an index of 1.45 and air with an index of 1.0.

1.2 Light, Waveguides and Photonic ICs

Waveguides are structures that can confine light within a certain area and transport it over a given distance. In almost all situations, this confinement is done by index guiding, where light is trapped in a core of material with a higher refractive index than the surrounding material [1]. Light is bundled in a spot around the centre, or a *mode*. Most optical fibres used in telecommunications support only one guided mode.

While optical fibres are made of glass, photonic ICs can be fabricated in a variety of materials. One of the most popular and mature commercial technologies today is *Silica on Silicon*. Just as in optical fi-

bres, the guiding material is glass, but it is arranged as a single glass layer on a flat surface in which waveguides with a higher refractive index are defined. The dimensions of the waveguide cores are comparable to those of single-mode fibres. An alternative material for PICs is semiconductor, like *Silicon* or more exotic $III-V$ semiconductors like *Gallium Arsenide* or *Indium Phosphide*. Semiconductors typically have a refractive index larger than 3, which is much higher than the index of glass. Silicon-on-insulator, a material system where the Silicon waveguide core is separated from the Silicon substrate by a layer of oxide, is a very attractive material for optical waveguides, as the top Silicon layer confines the light with its very high refractive index[53]. Unlike Silicon, $III-V$ semiconductors make it easy to integrate active functions into the chip, like the generation of light or conversion of light to electrical signals. However, recently substantial light emission has been demonstrated in Silicon [54], promising active functions in this material as well. In semiconductors, waveguides are often defined as ridges, where some material around the waveguide is etched away. In the vertical direction, light is confined by a layer with a higher index than the surrounding layer, while horizontally, the thicker ridge has a higher effective refractive index than the thinner surroundings.

In optical fibres, the difference in refractive index between core and cladding is typically very low. Because of this, the guiding is weak and the light in optical fibres can easily be made to escape from the core, by bending the fibre so the light can *miss the bend*. This starts to happen with a bend radius of $1cm$. The same happens to waveguide bends in photonic ICs made in Silica on Silicon. Because in semiconductor devices the refractive index contrast between waveguide core and cladding is higher, the effect will occur for smaller bend radii, typically of a few $100\mu m$. Because waveguides are an essential part of photonic ICs, and waveguide bends have to be large, this effect is a serious limitation to the reduction in size of photonic ICs.

A typical photonics component used today is the *arrayed waveguide grating*. While this component is very performing in splitting WDM signals into its component wavelength channels and recombining them, current implementations in low-contrast systems (like *Silica on Silicon*) measure multiple cm^2. This is mainly due to the restrictions on the bend radii. These large dimensions have a negative impact on the number of components that can be placed together on a chip, as well as on the fabrication yield.

This example shows that conventional waveguide technology poses limits to the integration of photonic functions onto a photonic IC. Therefore, ultra-compact waveguides are needed, that can confine light into smaller cores and guide it around sharper bends.

1.3 Nanophotonics

If, like in electronics, one wants to integrate many photonic functions onto one chip, the individual components should shrink considerably. Also, the chips themselves should shrink, in order to ease the fabrication yield and reduce the costs per component.

As electronics and photonics each have their own strengths, future photonic chips will probably combine both electronic and photonic functions. However, current photonic building blocks (e.g. waveguides, splitters, bends,...) are still orders of magnitude larger than their electronic counterparts (e.g. transistors, flip-flops,...), so large-scale integration of both electronics and photonics onto one chip is not cost-effective. Smaller photonic components are needed to bring the density of functional elements to the same level as that in current electronics. To accomplish this, structures with submicron features are needed, that can be fabricated with accuracies in the order of $10nm$. Therefore, this research field is often called *nanophotonics*.

For this reduction in size new photonic technology is needed to replace the conventional low-index-contrast waveguides with their large bend radius. *Photonic Wires* are a scaled-down version of these waveguides, with a very high refractive index contrast between core and cladding. Because of this high contrast, they can confine light in a very small cross section, and guide it through very sharp bends. Where conventional waveguides are typically much broader than $1\mu m$, photonic wires have dimensions of $500nm$ or less.

Photonic Crystals provide an alternative to fabricate very compact photonic components [4]. Unlike photonic wires, photonic crystals do not resemble a classical waveguide at all. Photonic crystals consist of a periodic stacking of wavelength-scale cells with a high contrast in refractive index in each cell. The most common geometries consist of a lattice of semiconductor rods in air, or a semiconductor slab with air holes drilled into it. Light is controlled by many strong reflections off these rods or holes, because of the strong refractive index contrast between the air and the semiconductor. Photonic crystals are further discussed in the next section.

In current planar waveguides the guided modes are not so different from the modes in single-mode optical fibres. Because of this, it is relatively easy to *squeeze* the light gently from the fibre to the waveguide and back using an *adiabatic taper*. However, with nanophotonic waveguides, the mismatch between the two modes is much larger. To couple the light from the fibre and back without excessive losses, either very long adiabatic tapers are needed, or a radically different approach is required. Here too, photonic crystals might provide a solution. This will be discussed briefly in chapter 3.

1.4 Photonic Crystals

Photonic crystals are periodic structures with periods of the order of the wavelength of the light and a very high refractive index contrast within each period [4, 3]. For telecommunications, where infrared light with wavelengths in the range $1.3\mu m - 1.6\mu m$ are used, the photonic crystal period is typically $0.5\mu m$ or less. The periodicity can extend in 1, 2 or 3 dimensions. Because of this high refractive index contrast, light will be scattered very strongly throughout the structure, and the scattered waves from each period can either add up or cancel out each other, depending on the wavelength of the light.

The most complex, and also the most interesting photonic crystals have a periodic nature in all 3 dimensions. For a well-chosen geometry and a unit cell with sufficiently high refractive index contrast the scattering from each cell can interfere in such a way that all light inside the crystal within a certain wavelength range is cancelled out, so no propagation is possible in the structure [2]. This wavelength range is called a *photonic bandgap* (PBG). The PBG makes it possible to control the flow of light in a structure. For instance, one could fashion a waveguide enclosed in a 3-D photonic crystal. This waveguide could have very sharp bands, because the photonic crystal forbids the propagation of light outside the waveguides. Light is therefore forced to follow the path of the waveguide. Similarly, one could trap light in a cavity enclosed by 3-D photonic crystal.

Although 3-D photonic crystals can control light in all directions, they are very difficult to fabricate for optical and infrared wavelengths. Their 3-D nature requires a very accurate stacking of elements with a high refractive index, and to introduce deliberate defects to serve as e.g. waveguides or cavities is far from obvious.

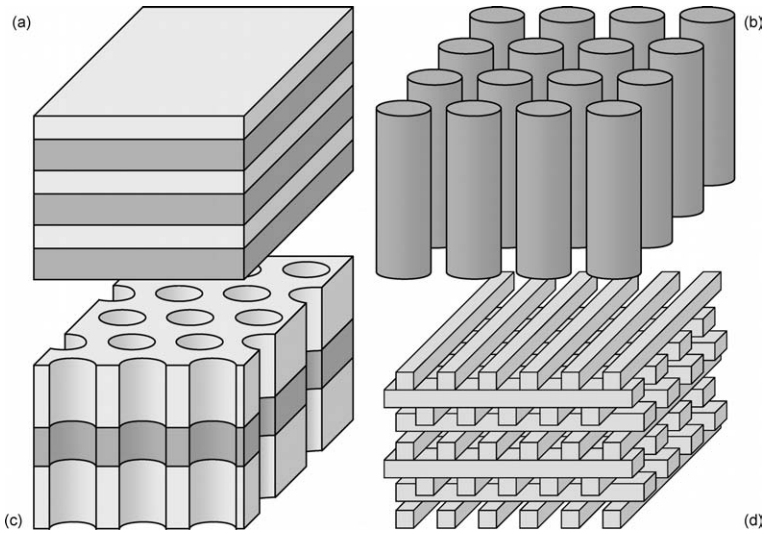

Figure 1.2: Some examples of photonic crystals. (a) A 1-D periodic Bragg stack. (b)A 2-D periodic square lattice of high-index rods in air. (c) A 2-D periodic photonic crystal slab consisting of air holes drilled into a slab waveguide. (d) a 3-D periodic stacking of high-index rods (the so-called *woodpile structure*)

An alternative to full 3-D photonic crystals is the combination of a conventional waveguide structure and a photonic crystal. Here, a 2-D photonic crystal is created by etching holes or rods in a semiconductor layer structure. In the horizontal direction the photonic crystal controls the flow of light, while in the vertical direction the light is confined in the layer with the higher refractive index. These *photonic crystal slabs* also provide 3-D control of light while they are much easier to fabricate. Figure 1.3 shows a single line defect waveguide in a photonic crystal slab. Horizontally, the light is confined by the photonic bandgap. However, unlike 3-D photonic crystals, the confinement in the vertical direction is not necessarily perfect. This can also be understood intuitively from figure 1.3. In the holes in the photonic crystal slab, the vertical contrast in refractive index is lost, so light is not confined. Because of this, light will scatter out toward the top or bottom and it cannot be recovered. These *out-of-plane scattering losses* are a serious restriction to the applicability of photonic crystal slabs. In chapter 4, we will discuss out-of-plane scattering in detail.

Because of the strong refractive index contrast, the subwavelength structures and the periodicity, photonic crystals are very hard to model.

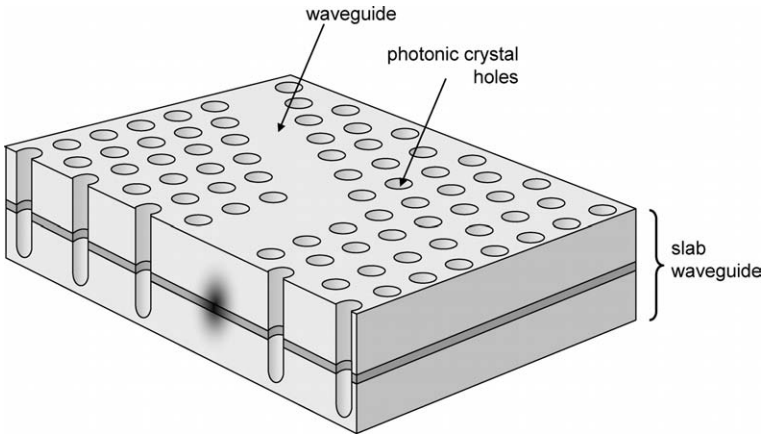

Figure 1.3: A Photonic crystal slab waveguide.The light is confined vertically by the layered structure of the slab waveguide. Horizontally, the photonic crystal prohibits light within the photonic bandgap to propagate through the lattice, effectively binding it to the defect of missing holes, and guiding it through the bend.

Various approximations that are valid for the simulation of conventional waveguides will not hold for photonic crystals. Simulation tools that can model photonic crystal structures rigourously in three dimensions are therefore very resource-intensive and require much calculation time. Therefore, the design of good components based on photonic crystals is in many cases a slow process of modelling, trial and error.

In chapter 2 we delve deeper into the properties of photonic crystals and photonic crystal slabs, while in chapter 4 we take a close look at the various loss mechanisms in photonic crystal slabs.

1.5 Nanoscale Fabrication

As already mentioned, nanophotonic components have dimensions of a few $100nm$ and require a fabrication accuracy of around $10nm$. Most conventional photonic components are fabricated with optical lithography, where a mask with the pattern of the component is projected onto a photosensitive film. The pattern on the film is then developed and transferred into the structure. While this is a very fast process, the problem with optical lithography is that the resolution is dependent on the illumination wavelength of the projection light source. For the optical

or near-UV wavelengths that are commonly used, this resolution falls dramatically short of what is needed for nanophotonics fabrication.

For nanophotonics fabrication for research purposes *e-beam lithography* is a very popular technology. Here, the pattern is written with a narrow electron beam in a vacuum system. This technique has a very high resolution, but because the patterns are written sequentially, it is very slow. Therefore, e-beam lithography is less suitable for mass-fabrication.

The obvious alternative is presented by the high-end electronics industry. State-of-the-art microprocessors contain transistors with dimensions around $100nm$, and the technology to make these structures is based on optical lithography. To attain the resolution needed for these very small structures, a shorter illumination wavelength in the deep UV is used. Typically, *deep UV lithography* is a form of optical lithography, but makes use of an excimer laser at $248nm$ or $193nm$, opposed to the $365nm$ wavelength used for the best conventional optical projection lithography. Like optical lithography, deep UV lithography projects a mask onto a *photoresist* layer, making it a very fast process. However, the process is optimised very much toward the fabrication of CMOS electronics, and it is not straightforward to transfer the fabrication know-how to the very different nanophotonic structures.

In this work, we have studied the potential of deep UV lithography for the fabrication of nanophotonic structures. For this purpose, we have set up a collaboration with *IMEC* in Leuven, partly in the context of the IST-PICCO project[1]. In chapter 5 we will explain the process in detail and discuss the various obstacles encountered.

1.6 Objectives

Nanophotonic waveguides in general, and photonic crystals in particular, are hard to understand, to model, to design and finally to (mass) fabricate. In this work we have tackled a number of these problems.

From the beginning we have focused on 2-D photonic crystal slabs, as these are the most suitable for integration into a PIC. At that time, photonic crystal slab waveguides had considerable propagation losses. One of our main objectives was therefore to understand the various loss mechanisms in photonic crystal slabs and find ways to reduce them to

[1] IST-PICCO project web site: *http://www.intec.ugent.be/picco*

reasonable levels. With this knowledge, we could then design better waveguides.

In parallel, we wanted to solve the fabrication issue. At the start of the research, fabrication was limited to slow e-beam writing. Access to the deep UV facilities in IMEC Leuven opened an avenue to acquire the necessary experience in CMOS processes to develop a suitable process for nanophotonic components.

1.7 Conclusion

In this introductory chapter we have tried to sketch the background and the rationale behind this work, as well as to introduce a few key topics. We modelled, designed and fabricated ultra-compact waveguide components, based on photonic crystals and photonic wires, for the use in photonic integrated circuits. We also studied the mass fabrication possibilities using deep UV lithography.

The remainder of this work is structured as follows. In the next chapter we will take a look at photonic crystals from a theoretical point of view, and focus on photonic crystal slabs. In chapter 3 we will then discuss some designs and design rules for photonic crystal slab waveguides. Photonic wire components are also treated in chapter 3, as well as couplers for interfacing between single-mode fibres and ultra-compact waveguides. Chapter 4 then goes deeper into the various loss mechanisms in nanophotonic waveguides. In chapter 5 the fabrication possibilities and limitations of deep UV lithography are discussed in detail. Finally, the fabricated components are measured and the measurements are correlated to the theoretical results in chapter 6. The main conclusions are put together in chapter 7.

Many parts of this work have been published in both peer-reviewed international journals and at international conferences. An overview of these publications is given in appendix B.

Chapter 2

Photonic Crystals

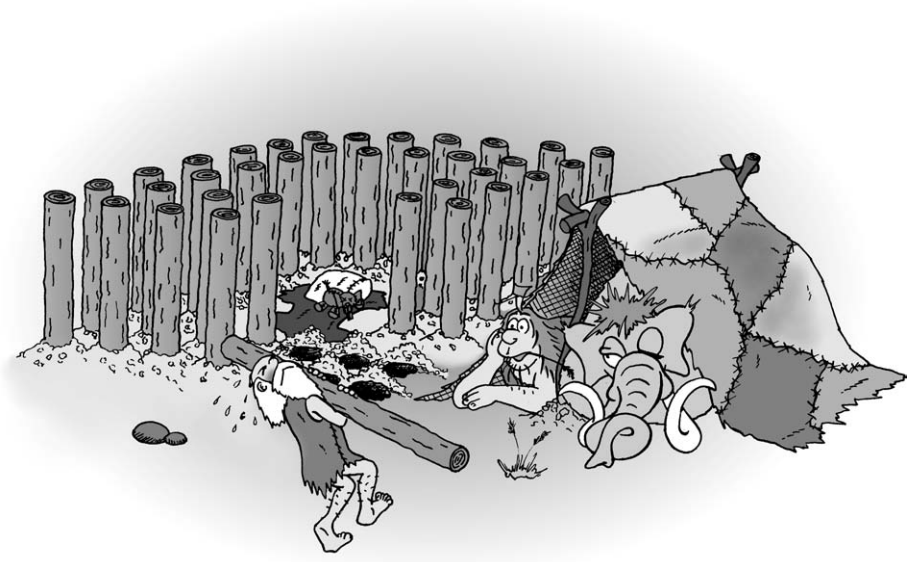

"Ug, I don't think this will keep out all the light."

Chapter 2

Photonic Crystals

In this chapter we will explain the fundamentals of the core structures of this work: *photonic crystals*. Photonic crystals are periodic electromagnetic structures. Therefore, we will briefly go over the Floquet-Bloch theory of periodic structures, and how they are applied to 1-D periodic media. We will then explain how 3-D periodic media can be used to completely control light. However, because 3-D structures are difficult to fabricate, we prefer to work with 2-D periodic structures fabricated in a slab waveguide, so-called photonic crystal slabs. To conclude this chapter, we will briefly discuss the various simulation tools that can be used to model photonic crystal structures.

2.1 What are Photonic Crystals?

The term *photonic crystal* is mainly used for 1-D, 2-D or 3-D periodic structures with *a period of the order of the wavelength of the light* and *a high refractive index contrast in each unit cell*. The concept was first introduced by Eli Yablonovitch in 1987 as a method of controlling spontaneous emission in a material [2]. In a periodic medium, electromagnetic waves scattered within each period can either add up or cancel out. Because of this *interference*, the structure can become transparent or opaque, depending on the direction and the wavelength of the electromagnetic waves. For some geometries and a sufficiently high refractive index contrast a multi-dimensional photonic crystal can have a wavelength (or corresponding frequency) range where electromagnetic

radiation cannot propagate through the structure in any direction. This wavelength range is called a *photonic bandgap*[1] or *PBG* [3].

Photonic crystals come in many forms. The most obvious classification is by the dimensionality of the periodicity. Figure 1.2 on page 9 shows 4 kinds of photonic crystals with different periodicity. The most simple, 1-D photonic crystals have been known and used for a long time in the form of Bragg mirrors. These layered structures have a PBG for light incident perpendicular to the stack. As the angle of incidence moves away from the vertical direction, the PBG wavelength range decreases and ultimately closes. Both the width of the PBG and the maximum angle for which there is such a forbidden wavelength range is determined by the refractive index contrast between the layers; the higher the contrast, the broader the PBG [4]. In the next section, we will briefly go into the theory of these 1-D photonic crystals, as we will explain the behaviour of light in periodic media.

When the periodicity is extended into other directions, the picture becomes more complicated. For each direction, the relationship between reflected electromagnetic waves is different, and the wavelength range for which the waves cancel out is shifted. Therefore, to obtain a full photonic bandgap, these wavelength ranges should overlap for all directions. This is only possible when the refractive index contrast within the structure is sufficiently high and the structure has a high symmetry.

2.2 Light in Periodic Media

In order to get a more profound understanding of the behaviour of photonic crystals, we will first study electromagnetic waves in a 1-D periodic medium. We will then extend these insights to more dimensions, and apply them to photonic crystals.

[1] The term *photonic bandgap* is reminiscent to the *"electronic"bandgap* found in very pure crystalline materials. There, a bandgap denotes a range of energies where no electrons can be found inside the bulk of the (defect-free) crystal.

2.2.1 Light in Homogeneous Media

Light is a form of electromagnetic radiation, governed by *Maxwell's equations*:

$$\begin{aligned} \nabla \times \mathbf{E} &= -\frac{\partial \mathbf{B}}{\partial t} \\ \nabla \times \mathbf{H} &= \frac{\partial \mathbf{D}}{\partial t} + \mathbf{J} \\ \nabla \cdot \mathbf{D} &= \rho \\ \nabla \cdot \mathbf{B} &= 0. \end{aligned} \tag{2.1}$$

In a macroscopic world, where we disregard the quantum nature of matter, the electromagnetic fields are linked by constitutive relations:

$$\begin{aligned} \mathbf{D} &= \epsilon \mathbf{E} = \epsilon_r \epsilon_0 \mathbf{E} \\ \mathbf{B} &= \mu \mathbf{H} = \mu_r \mu_0 \mathbf{H}. \end{aligned} \tag{2.2}$$

In free space, or in a linear, isotropic and homogeneous (*LIH*) medium, ϵ_r and μ_r are a scalar function of the frequency. For our convenience, we also assume that there is no material dispersion, hence no frequency-dependence, so ϵ_r and μ_r are constant. When we now insert the constitutive relationships for these materials into the Maxwell equations, and set all currents $\mathbf{J}$ to zero, we can rewrite the equations in the following form:

$$\begin{aligned} \nabla^2 \mathbf{E}(\mathbf{r}, t) - \epsilon\mu \frac{\partial^2 \mathbf{E}(\mathbf{r}, t)}{\partial t^2} &= 0 \\ \nabla^2 \mathbf{B}(\mathbf{r}, t) - \epsilon\mu \frac{\partial^2 \mathbf{B}(\mathbf{r}, t)}{\partial t^2} &= 0. \end{aligned} \tag{2.3}$$

The solutions of these wave equations are plane waves, travelling at the speed v, which is determined by the dielectric constant ϵ and magnetic permittivity μ of the material:

$$\begin{aligned} v &= \frac{1}{\sqrt{\epsilon\mu}} \\ &= \frac{c}{\sqrt{\epsilon_r \mu_r}} \\ &= \frac{c}{n}, \end{aligned} \tag{2.4}$$

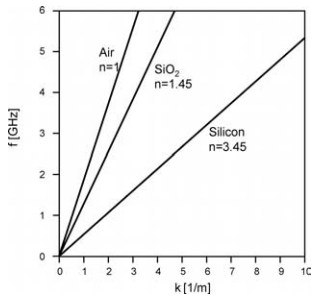

Figure 2.1: The dispersion relation for LIH media with refractive index n of 1 (Air), 1.45 (Silica or glass) and 3.45 (Silicon).

with c the speed of light in vacuum, and n the *index of refraction*. We can thus rewrite the equations as

$$\begin{aligned} \frac{c^2}{n^2}\nabla^2\mathbf{E}\left(\mathbf{r},t\right) - \frac{\partial^2\mathbf{E}\left(\mathbf{r},t\right)}{\partial t^2} &= 0 \\ \frac{c^2}{n^2}\nabla^2\mathbf{B}\left(\mathbf{r},t\right) - \frac{\partial^2\mathbf{B}\left(\mathbf{r},t\right)}{\partial t^2} &= 0. \end{aligned} \tag{2.5}$$

The general solution for the electric field $\mathbf{E}$ on time t and in position $\mathbf{r}$ is a superposition of plane waves of the form

$$\mathbf{E}\left(\mathbf{r},t\right) = \mathrm{Re}\left[\mathbf{E}_0 e^{j\mathbf{k}\cdot\mathbf{r}} e^{-j\omega t}\right], \tag{2.6}$$

where $\mathbf{E}_0$ is complex, describing both the amplitude and the phase of the plane wave. f is the frequency of the plane wave, and is related to ω as

$$\omega = 2\pi f. \tag{2.7}$$

$\mathbf{k}$ is the wave vector, indicating the propagation direction of the wave. The magnitude of $\mathbf{k}$, k is linearly dependent on f, or ω:

$$\frac{\omega}{\|\mathbf{k}\|} = \frac{\omega}{k} = \frac{c}{n}. \tag{2.8}$$

The relation between $\mathbf{k}$ and f or ω is called the *dispersion relation* and is of course dependent on the medium through the refractive index n. Figure 2.1 shows the dispersion relation for a plane wave propagating in the x direction in vacuum ($n = 1$), in glass ($n = 1.45$) and in Silicon ($n = 3.45$) in a $f\|k$ diagram. We see that the slope of the dispersion curve is lower as the refractive index of the material increases.

$\mathbf{k}$ and f are directly related to the wavelength λ in the medium:

$$\lambda = \frac{2\pi}{k} = \frac{c}{f} = \frac{\lambda_0}{n}, \tag{2.9}$$

with λ_0 the wavelength in vacuum.

In an LIH material, monochromatic plane waves propagate at the speed c/n, or ω/k. However, when light of a certain wavelength is modulated by a signal, the resulting beam will not be monochromatic, but will use a certain wavelength band around the carrier wavelength. Here, the modulation (or amplitude envelope) and the phase fronts of the carrier wave do not necessarily travel at the same speed. While the phase fronts travel with a phase velocity $v_p = \omega/k$, the group velocity v_g of the modulation (or the *information*) is

$$v_g = \frac{\partial \omega}{\partial k} = 2\pi \frac{\partial f}{\partial k}. \tag{2.10}$$

In LIH materials, which have a linear dispersion relation, $v_p = v_g$. However, as we will see in the next section, this changes when the material is no longer homogeneous.

However, as long as the materials are linear, i.e. ϵ and μ are constant and the same for all frequencies, we can write the general solution of the Maxwell equations as a sum, or integral of monochromatic solutions. This allows us to treat many problems in the frequency domain, where we assume monochromatic solutions in advance [1].

Maxwell's equations contain no fixed scale. As long as the material is macroscopic compared to the wavelength of the radiation, the solution is valid for all frequencies. This can also be seen on the straight dispersion curve, which looks the same on any scale.

However, how we perceive electromagnetic radiation depends on the frequency or wavelength. Figure 2.2 shows the electromagnetic spectrum with the low-frequency radio waves, the range which we perceive as *light*, including infrared and ultraviolet, and ionising radiation, like X-rays and Gamma rays, at the high-frequency end. To us, the region near the visible light is the most interesting. Our eyes are sensitive for a narrow wavelength range between red near $700nm$ and violet near $400nm$, the *visible spectrum*. For telecommunication applications, the near infrared is more interesting: optical fibres from high-purity glass have a very low absorption around $1.55\mu m$ and $1.3\mu m$. With deep ultraviolet light of wavelengths of $248nm$ and $193nm$ we can define and fabricate very small features, as we will discuss in chapter 5.

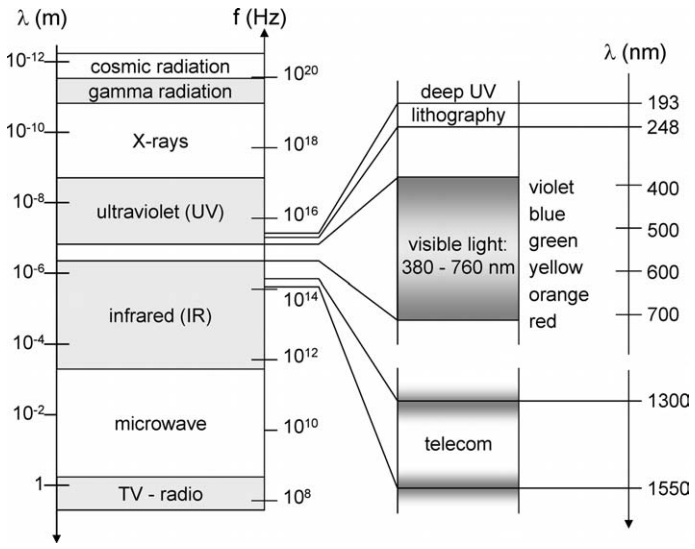

Figure 2.2: The electromagnetic spectrum: Note that the frequency scale is reversed with respect to the wavelength scale.

Even though radiation for different frequencies seems to behave very differently, this is only because we use ourselves as a scale reference, and because materials react differently at different wavelengths. However, if the material properties were wavelength-independent, the equations would scale throughout the spectrum. When we will discuss periodic structures in the next paragraphs, the results derived are equally valid for radio and microwaves as for visible light. Only, because we develop applications for optical telecommunication, we will concentrate on the wavelength range near $1.55\mu m$.

2.2.2 1-D Periodic Media

Plane waves describe the propagation of light in a uniform medium. However, photonic crystals are far from uniform, and the wave equation for $\mathbf{E}$ will change to

$$\frac{c^2}{n^2(\mathbf{r})}\nabla^2\mathbf{E}(\mathbf{r},t) - \frac{\partial^2\mathbf{E}(\mathbf{r},t)}{\partial t^2} = 0 \tag{2.11}$$

To simplify the problem, let's consider a one-dimensional medium, with waves propagating along the z-axis. This simplifies the wave

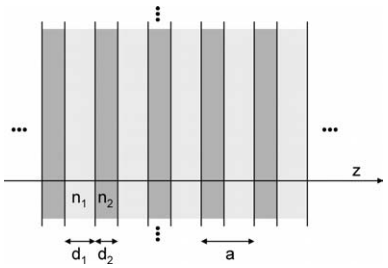

Figure 2.3: A 1-D periodic medium consisting of a stack of piecewise constant refractive index materials

equation:

$$\frac{c^2}{n^2(z)}\frac{\partial^2 \mathbf{E}(z,t)}{\partial z^2} - \frac{\partial^2 \mathbf{E}(z,t)}{\partial t^2} = 0 \tag{2.12}$$

.

In a homogeneous medium, where $n(z)$ is constant, the solution of the wave equation will become a plane wave travelling in the $+z$ or $-z$ direction:

$$\mathbf{E}(z,t) = \mathrm{Re}\left[\mathbf{E}_0 e^{jk_z(\omega)z} e^{-j\omega t}\right]. \tag{2.13}$$

Let us now consider an infinite 1-D periodic medium with period a, like the layer structure in figure 2.3, $n(z)$ becomes a periodic function of z. We consider the problem to be uniform in the other directions and only study waves travelling along the z-axis. According to the *Floquet-Bloch theorem* the solution of the wave equation is of the form

$$\mathbf{E}(z,t) = \mathrm{Re}\left[\mathbf{u}_k(z) e^{jk_z(\omega)z} e^{-j\omega t}\right], \tag{2.14}$$

where $\mathbf{u}_k(z)$ is periodic with the same period a as $n(z)$. For linear materials, the general solution consists of a linear superposition of these *Floquet-Bloch waves*. Note that we have now introduced a preferential scale, and that the solutions of the wave equation do not scale arbitrarily, but scale with the period a.

For this periodic medium, k_z is no longer linearly dependent on ω or f, and for each frequency there are an infinite number of solutions for k_z that are coupled to one another by the periodic lattice. In figure 2.4 we can see the dispersion relation change as we increase the refractive index contrast in the layer structure. We can see that the dispersion relation is periodic in k_z with a period $2\pi/a$, and symmetric around $k = 0$. Because of the lattice periodicity all points in the k-space are coupled to a point in the interval $[-\pi/a, \pi/a]$, called the *first Brillouin*

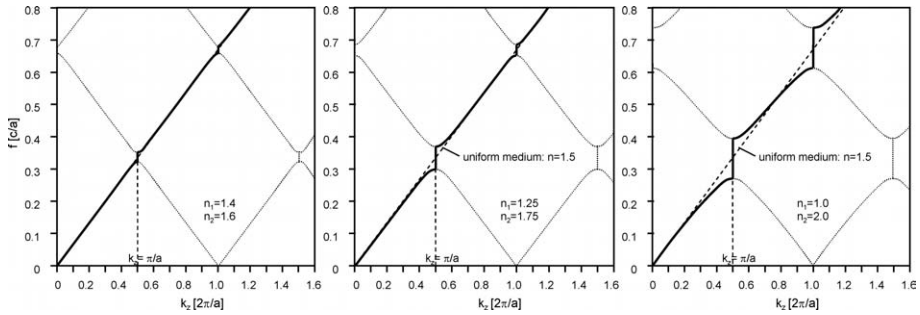

Figure 2.4: The dispersion relation for the 1-D periodic medium of figure 2.3 with (a) $n_1 = 1.4$ and $n_2 = 1.6$, (b) $n_1 = 1.25$ and $n_2 = 1.75$, (c) $n_1 = 1$ and $n_2 = 2$. $n_1 d_1 = n_2 d_2$. The linear dispersion relation of a uniform medium with $n = 1.5$ is also plotted.

zone. Because of the reciprocity of the Maxwell equations in the absence of an external magnetic field, the dispersion relation in the $-z$ direction is the same as in the $+z$ direction. It therefore suffices to study the dispersion relation for $k \in [0, \pi/a]$, the *reduced first Brillouin zone*, or just the *reduced Brillouin zone*.

The branches with a positive (upward) slope represent forward-propagating waves coupled to the lattice. The branches with a downward slope are the result of the light propagating in the $-z$ direction, but coupled by a lattice vector to the first reduced Brillouin zone. As the group velocity is given by the slope of the dispersion relation, written in equation 2.10, these branches have a negative group velocity. Therefore, the phase fronts will propagate in the opposite direction as the modulation envelope of the light.

For low frequencies, the material behaves more or less like a homogeneous medium, because the period a is much smaller than the wavelength of the light, and index variations are averaged out. However, for higher frequencies, k_z no longer behaves linearly, but splits in 2 branches that flatten out. In the area with low slope, the group velocity v_g of the light waves becomes very small, and the light comes to a halt, resulting in a standing wave at the point where the dispersion is horizontal. In the frequency interval between the two branches, there is no real solution for k_z, but k_z becomes complex. The imaginary part translates into an attenuation in the propagation direction. In an infinite periodic medium, this means no propagating mode is possible in that frequency range. The frequency or wavelength range in which this happens is called the *photonic bandgap* or *PBG*, and gets larger for higher

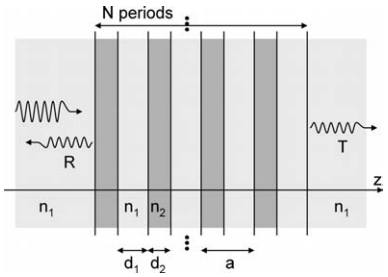

Figure 2.5: A finite 1-D Bragg stack with N periods. The incident plane wave from the homogeneous medium is partly reflected and partly transmitted.

index contrast in the periodic structure. The frequency domain is split up by the dispersion relation into photonic bandgaps or *stop bands* and transmission or *pass bands*. Because of this, we call the dispersion plot of a periodic structure a *band diagram*.

In a finite periodic medium, the Floquet-Bloch theorem is no longer valid, but if the medium is large enough, it still provides a very good approximation. Let's look at the case of a 1-D periodic stack of finite thickness, as illustrated in figure 2.5. If we let a plane wave from the left uniform medium fall onto the stack, we can calculate the reflection and the transmission of the light as a function of wavelength or frequency. In figure 2.7 we see a decrease in transmission and increase in reflection in the wavelength range near $\lambda = 2p\,(n_1 d_1 + n_2 d_2)$. with p a whole number. This is called the *Bragg condition* and the centre wavelength is called the *Bragg wavelength*. At this point, the waves reflected by each even interface are in phase, and interfere constructively. The same is true for the reflected waves at each odd interface. Additionally, in the case where $n_1 d_1 = n_2 d_2$, the even and odd reflections are in phase. This is called the *Mie condition*. This is illustrated in figure 2.6.

Slightly away from the Bragg condition, this interference is not perfect, but it is still enough if the individual reflections at each interface are strong enough. This can be done by using a layer stack with high index contrast. This makes the photonic bandgap wider, as we have also seen in figure 2.4.

As we increase the number of periods N, the reflection in the photonic bandgap becomes stronger. If we add enough periods, we can thus make an arbitrarily strong reflector. These *Bragg mirrors* have a number of applications, especially where absorbing metals are to be avoided. Also, the stronger the index contrast in the periodic structure,

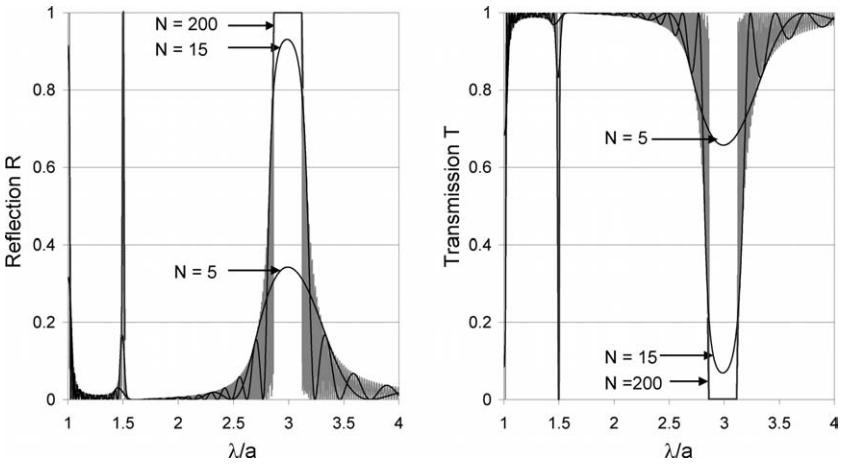

Figure 2.6: Transmission and reflection spectrum of the finite 1-D Bragg stack from figure 2.5 for various thicknesses N of the stack. $n_1 = 1.4$, $n_2 = 1.6$ and $n_1 d_1 = n_2 d_2$.

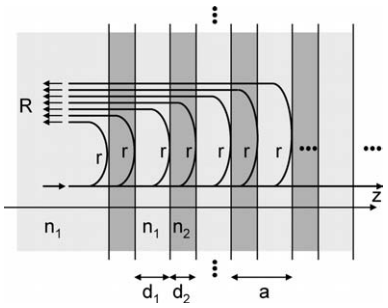

Figure 2.7: The Bragg condition in a 1-D periodic medium: when the reflections from the various interfaces are in phase, they interfere constructively. This is the case when $\lambda = 2p\,(n_1 d_1 + n_2 d_2)$. With sufficient periods, the reflection R goes to 1.

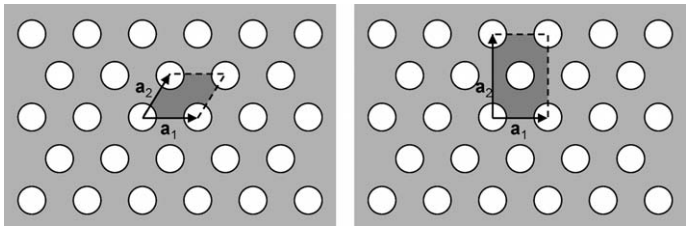

Figure 2.8: A 2-D periodic structure, with 2 sets of base lattice vectors and unit cells. The structure is uniform in the out-of-plane direction.

the more effective the Bragg mirror becomes, both in wavelength range and in reflectivity.

In photonics applications, Bragg mirrors are often implemented as layered stacks of monocrystalline semiconductors. Because the fabrication of such layer structures is subject to a number of limitations, the range of thicknesses or index contrasts will be limited.

2.2.3 Periodicity in Multiple Dimensions

Now the theory of 1-D periodic structures can be extended to two and three dimensions. For 1-D periodic structures, the periodicity can be expressed with a single base lattice vector $\mathbf{a}$ and a *unit cell*. The base lattice vector $\mathbf{a}$ has a length equal to the lattice period a and the direction of the periodicity. When we add dimensions to the periodicity, we can assign additional base lattice vectors to the structure, which we can name $\mathbf{a_1}$, $\mathbf{a_2}$, ... The structure is then periodically modulated along these directions. Note that it is possible to choose a number of equally valid sets of base lattice vectors for a single periodic structure, as long as the vectors span the entire lattice. This is illustrated in figure 2.8. With the choice of lattice vectors comes a unit cell, which is translated along any integer linear combination of the base lattice vectors.

In a 2-D or 3-D periodic structure, $n(\mathbf{r})$ is a periodic function with lattice vectors $\mathbf{a_1}$ and $\mathbf{a_2}$ (and $\mathbf{a_3}$ in three dimensions. Let's consider 2-D periodicity for the remainder of this section; the theory is easily extended to three dimensions.) In this case, the Floquet-Bloch theorem states that the solution of wave equation is of the form:

$$\mathbf{E}(\mathbf{r},t) = \mathrm{Re}\left[\mathbf{u_k}(\mathbf{r})e^{j\mathbf{k}(\omega)\cdot\mathbf{r}}e^{-j\omega t}\right]. \tag{2.15}$$

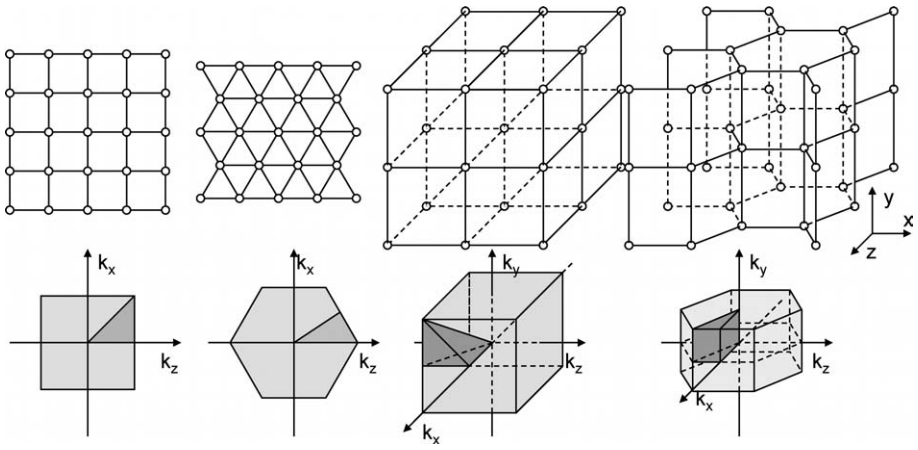

Figure 2.9: 2-D and 3-D periodic structures with their first Brillouin zone. Left to right: A 2-D square lattice, a triangular lattice, a simple cubic lattice, a hexagonal lattice. Bottom row: k-space with the first Brillouin zone and the reduced Brillouin zone.

The function $\mathbf{u_k}(\mathbf{r})$ is periodic with the same periodicity as $n(\mathbf{r})$. The length of the propagation vector $\mathbf{k}$ is now no longer linear in ω, and neither is it the same in all directions.

Just as in 1-D, the points in 2-D or 3-D k-space are coupled to the first Brillouin zone by the lattice periodicity. With lattices that have a rotational symmetry, the dispersion relation has the same symmetry, and we can again limit our study to the reduced Brillouin zone. Figure 2.9 shows some lattices with their respective first Brillouin zones.

2.2.4 2-D Periodic Structures

As an example, we will look at the case of a high-index material $n_2 = 3$ with a 2-D triangular lattice of round holes with index $n_1 = 1$ in it. The lattice extends in the xz-plane and is assumed uniform in the out-of-plane y-direction. This structure is illustrated in figure 2.10. We have chosen 2 lattice vectors $\mathbf{a_1}$ and $\mathbf{a_2}$. The first Brillouin zone has a hexagonal shape. Because of the sixfold symmetry of the lattice, the reduced Brillouin zone is a right triangle. We call the centre point of the reciprocal space Γ, the point in the direction of the nearest hole is called K, and the remaining point is M.

As we assume the 2-D structure is invariant in the third direction y, and study only waves with no propagation component in that di-

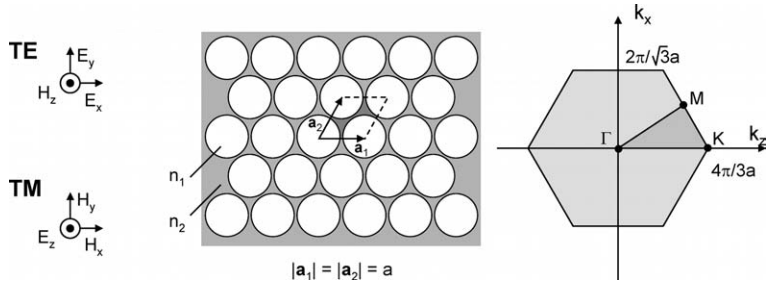

Figure 2.10: 2-D Triangular lattice of air holes in a high-index material. Left to right: the structure with the unit cell demarcated, the triangular lattice and the first and reduced Brillouin zone.

rection ($k_y = 0$), we can split up the full vectorial problem into two uncoupled polarisations: The *TE polarisation* (Transverse Electric field) with the electric field in the plane of the structure, and the *TM polarisation* (Transverse Magnetic field) with the electric field perpendicular to the structure. For the magnetic field, the roles are reversed[2]. The orientation of the **E** and **H** fields are also indicated in figure 2.10.

If we now plot the dispersion relation of this periodic structure for **k** in the reduced Brillouin zone defined by Γ, K and M we get the surface plots of figure 2.11 for the TE and the TM polarisation. We see that the dispersion relation is again split up, but the frequency at which this happens is different for the two polarisations and for the various directions of **k**. For each direction, there is a photonic bandgap. However, the bandgaps for all directions and polarisations not necessarily overlap.

A band diagram plotted as a surface plot is hard to study quantitatively. Therefore, instead of studying the dispersion relation in the

[2]There is some controversy as for the use of TE and TM as we use it here in the context of photonic crystals. In some works TE and TM are substituted by s and p , or E and H, in order not to confuse them with the use of TE and TM in slab and ridge waveguides. However, as TE and TM are the most widely used, we have chosen to do the same.
This choice of TE and TM is seemingly contradictory to the definition used in layered media. There, TE is the situation where the electric field is parallel to the interfaces between materials. This convention was conserved in waveguide theory, where (quasi-)TE designated the situation where **E** is parallel to the horizontal interfaces of the slab waveguide, defined by the xz-plane. This directly explains the convention in 2-D photonic crystals. Because many 2-D photonic crystals are actually created by drilling holes (or erecting pillars) in a slab waveguide in the xz-plane, we stick to the convention that the TE polarisation is the one with the electric field in the xz-plane.

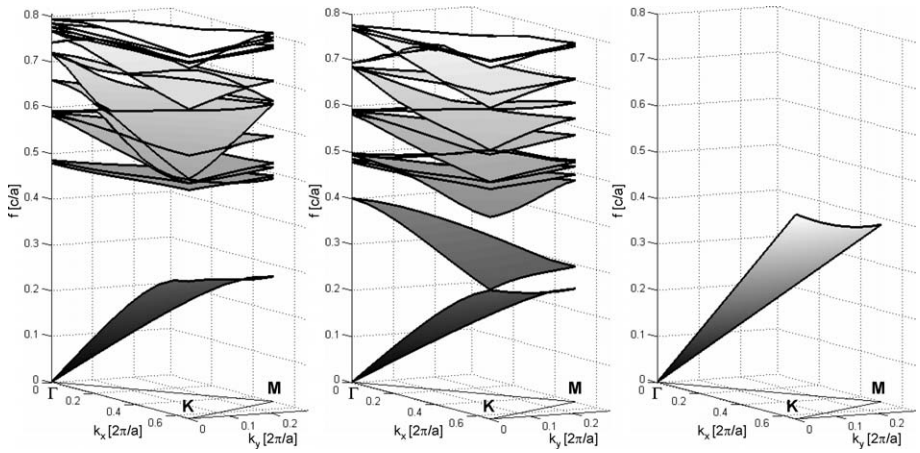

Figure 2.11: Dispersion relation for the structure in figure 2.10. Left to right: TE polarisation, TM polarisation, uniform medium with $n = 1.5$.

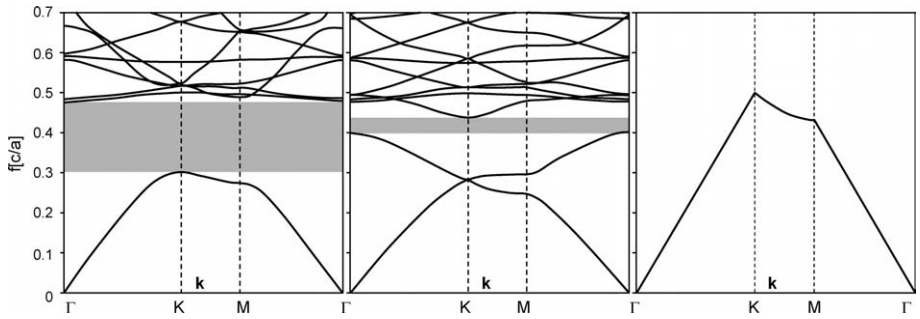

Figure 2.12: Band diagram for the structure in figure 2.10. Left to right: TE polarisation, TM polarisation, uniform medium with $n = 1.5$. $r/a = 0.3$

whole reduced Brillouin zone, we will only look at it along the edges between Γ, K and M. In this way, we can get a clear band diagram, as plotted in 2.12.

For the TE polarisation, we can see that there exists a large frequency range where no propagating modes are allowed in the structure. This range is called an *omnidirectional* photonic bandgap. For TM, this frequency range is much smaller. The region where it overlaps with the PBG for the TE polarisation is called a *full photonic bandgap*. In this frequency or wavelength range, no light is allowed in the infinite periodic structure.

Now while in one dimension, the splitting of the dispersion relation into bands automatically led to a photonic bandgap, in 2-D and 3-D the

emergence of a full PBG is not guaranteed. Because the structure is not identical in all directions, the bandgaps of all directions can only overlap if they are wide enough, and if their centre frequencies are not offset too much from each other. The first condition, making the bandgap large enough, can be brought about by increasing the refractive index contrast in the unit cell. The 3-to-1 index contrast we used in our example already gives a large PBG for the TE polarisation. To bring the centre frequencies of the PBGs close together one has to make the structure more symmetric, in order to obtain a first Brillouin zone that resembles a circle as much as possible. The triangular lattice with its hexagonal first Brillouin zone already has the optimal symmetry.

In figure 2.12 we could see that TE had a much larger PBG than TM. Because of the different field orientations, the behaviour of TE and TM light in a 2-D periodic structure is very different. Without going into details, we can say that TE will scatter more strongly (and therefore have a larger bandgap) in an interconnected network of high-index material. This is definitely the case in the triangular lattice of air holes that we used as an example. In the TM polarisation, on the other hand, light reacts more strongly to lumped nodes of high-index material, like high-index pillars in an air surrounding, as shown in the right part of figure 2.13. In our previous example, the material between 3 holes could be regarded as a lumped node, and therefore we still have a PBG for TM, albeit somewhat smaller than the one for TE [4].

In this work, we will mainly focus on photonic crystals for the TE polarisation, and most structures will consist of a triangular lattice of holes in a high-index material, mostly semiconductor. Therefore, unless otherwise stated, if we mention *photonic bandgap* for a 2-D structure, it will imply an *omnidirectional* bandgap for the TE polarisation only, unless otherwise stated.

2.2.5 3-D Periodic Structures

When we extend the periodicity to three dimensions, polarisations are no longer independent, and we have to go back to a full vectorial treatment. We will not go deeply into 3-D photonic crystals in this work, but will limit us to some examples.

Let's take a look at the structures illustrated in figure 2.14. These structures are packings of spheres, but in the top structures the spheres are composed of high index material in air, while in the bottom structures the spheres are cut out of a block of high index material. The cubic

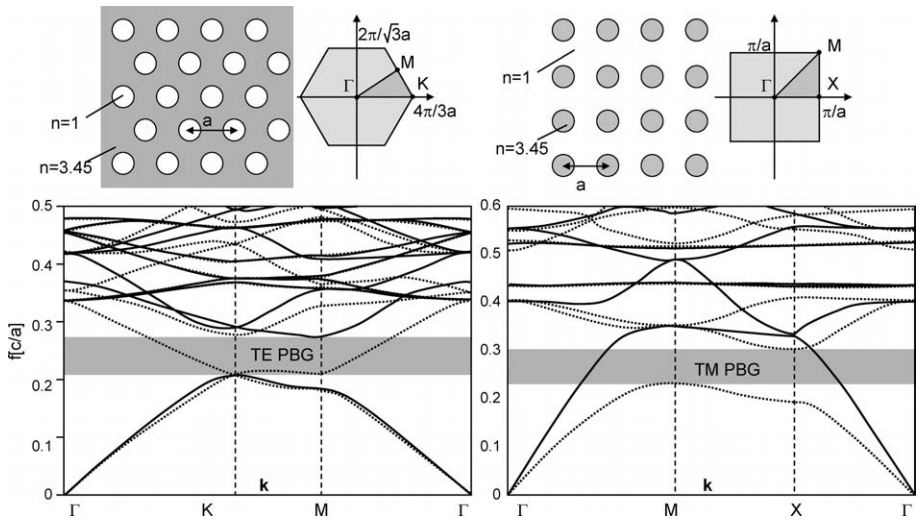

Figure 2.13: TE vs. TM: In some structures, there is a bandgap for the TE polarisation, while other structures only support a TM photonic bandgap. In the band diagrams, TE is plotted as a continuous line, while TM is the dotted line. Left: The band diagram for a triangular lattice of small air holes in a material with $n = 3$. This structure has a bandgap for TE, but not for TM. Right: A triangular lattice of pillars with index $n = 3$ in air. This time, TM has a PBG, but not TE.

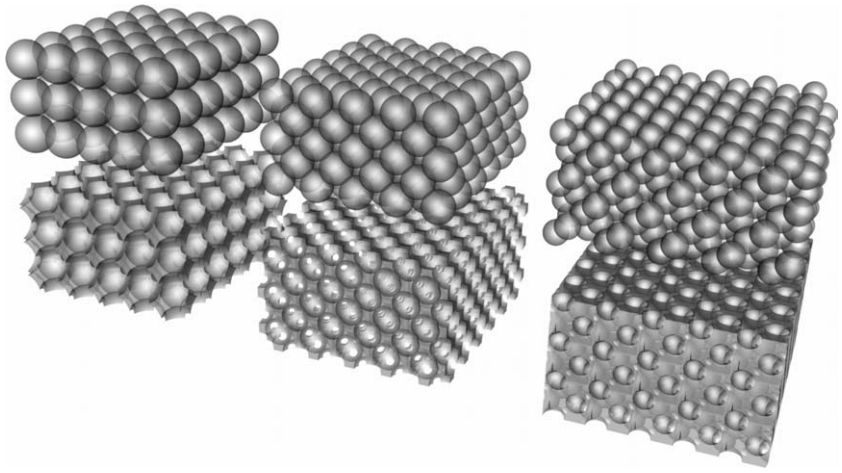

Figure 2.14: 3-D periodic structures consisting of a stacking of spheres. In the top row high-index spheres are stacked in a low-index background, while the bottom structures the spheres are cut out from a high-index material. Left to right: Simple cubic lattice, face-centered cubic closest packing, diamond lattice

lattice on the left has little symmetry. The *face-centered cubic* (FCC) lattice in the middle, which is a closest packing lattice, has much more symmetry. The Brillouin zone is illustrated in figure 2.15, and closely resembles a sphere. The right part of figure 2.14 is the diamond lattice, and it has the same Brillouin zone as the face centered cubic lattice.

Just as with 2-D structures, we can trace the dispersion relation around the reduced Brillouin zone. Figure 2.16 plots the band diagram for the FCC lattices and diamond lattices in figure 2.14. It is obvious that the diamond lattice is much better for the creation of a photonic bandgap than the closest packing, even though it has a somewhat lower symmetry. Both the stacking of high-index spheres and the inverted structure have a broad PBG.

These 3-D photonic crystals can be made by stacking spheres. In the macroscopic world, this can be done mechanically. For optical wavelengths, the spheres have submicron dimensions, and the stacking can be done with self-assembling colloidally suspended spheres. These structures are called *synthetic opals*, because they resemble natural opals in both structure and opalescence. When the space between spheres is filled up with a high-index material and the spheres themselves are removed, the resulting structure, called an *inverted opal*, will have a full photonic bandgap. Both structures are illustrated in figure 2.17 [55].

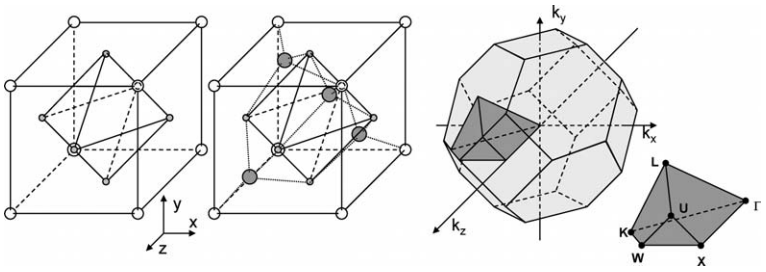

Figure 2.15: 3-D lattices with their Brillouin zone. Left to right: face-centered cubic lattice, diamond lattice, the first Brillouin zone and the reduced Brillouin zone.

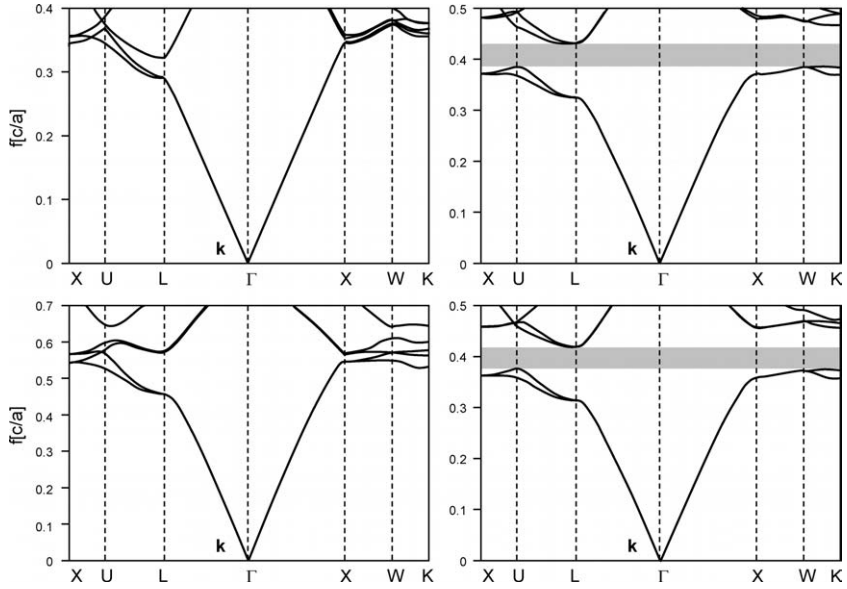

Figure 2.16: Band diagrams of the middle and right structures from figure 2.14, with $n_1 = 1$ and $n_2 = 3.45$. Top: Opals. Bottom: Inverted opals. Left: Face-centered cubic lattice. Right: Diamond lattice.

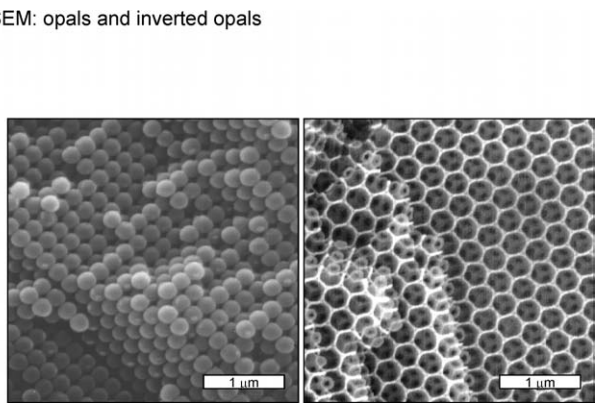

Figure 2.17: Opals and inverted opals. Courtesy of the Optoelectronics Research Group of Glasgow University [55].

While the diamond lattice has much better properties than the ordinary FCC lattice, it is not as easy to fabricate. Two examples are given in figure 2.18. The first, called *Yablonovite* was proposed by Eli Yablonovitch [56], and consists of a high index material with cylindrical holes drilled in three directions into it. This structure has a lattice that is very close to a full diamond lattice, but for a full diamond lattice, holes in another three directions are required. The structure was first made for microwave frequencies. At optical wavelengths, this structure would require etching deeply in three directions, which is a nontrivial procedure.

Another example is the woodpile structure, also illustrated in figure 2.18. This structure does have a full diamond lattice, and can be fabricated for optical wavelength by a sequence of lithography and etching steps [57].

2.2.6 Conclusion

We have shown that when we study light in a periodic medium, the dispersion relation is changed by the interaction between the light and the periodic lattice, creating bands where transmission is allowed, and frequency bands where no light can travel through the structure, the *photonic bandgaps* (PBG). In one dimension, we can thus create Bragg mirrors by periodically stacking layers with different refractive indices. The stronger the refractive index contrast, the broader the photonic bandgap becomes.

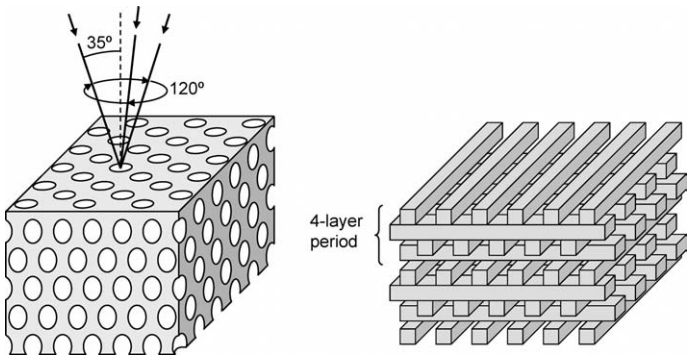

Figure 2.18: Photonic crystals with a diamond or diamond-like lattice. Left: Yablonovite, fabricated by drilling holes in three directions into a high-index material. Right: The woodpile structure, consisting of four stacked layers of rods.

In more dimensions, the frequency ranges where no propagation is allowed differ in the various propagation directions. In order to get a full PBG, i.e a photonic bandgap in all directions for all polarisations, we need a lattice structure with high symmetry and a sufficiently high refractive index contrast. Typically, for a 2-D periodic structure we need an index contrast of 2 to 1, while in 3-D a higher index contrast of 3 to 1 is needed. For optical wavelengths, the most suitable materials for this purpose are semiconductors, which can have refractive indices of 2.7 and upward.

3-D photonic crystals have the best properties to control light, because of their omnidirectional PBG. However, they are not trivial to fabricate, let alone integrate in planar photonic structures. Therefore, 2-D photonic crystals show the most promise. As we will explain in section 2.5, when we combine the 2-D periodicity with a slab waveguide, we can keep the light in the plane of the photonic crystal.

2.3 Defects in PhCs

In the previous section, we have discussed the properties of perfect, infinite periodic structures. While the infinity can be approached by making the structure sufficiently large, the applications of the perfectly periodic structure are limited. We have shown the use of a 1-D Bragg reflector, and a multi-dimensional photonic crystal can be used as an om-

nidirectional reflector[3]. Another use of the periodic structure is to use the very flat dispersion relations near the edge of the photonic bandgap (the *band edge*). At that point, the waves are slowed down, an effect that can be used for optical delay lines, or in situations where distributed feedback should be very strong, like in DFB-lasers. Because photonic crystals are so dispersive in this region, the refraction of light of nearby wavelengths is very different. This dispersive effect can be used to split an incoming beam into its constituent wavelengths.

While these applications have their use, the possibilities of photonic crystals increase dramatically when we introduce defects. Just like the photonic crystal can be periodic in 1, 2 or 3 dimensions, so can the defects be either point-like, or localised along a line or a plane. Figure 2.19 illustrates some defects in a 2-D triangular lattice. Introducing a defect in a photonic crystal breaks the translation symmetry of the periodic structure. Therefore, a new state can be created in the photonic bandgap, allowing light to exist in the neighbourhood of the defect. However, because the defect is surrounded by photonic crystal, the light cannot propagate away from the defect. In this way, we can create a cavity from which light cannot escape, or a line defect where light can only propagate up and down the direction of the defect. For a finite photonic crystal, the transition surface between the periodic structure and the uniform medium is also a defect, where under some conditions bound states can exist.

In a cavity in a 3-D periodic photonic crystal, light can be perfectly confined. A light emitter in that cavity can therefore only generate light in the bound state of the cavity. If we now provide a single escape route from that cavity, e.g. by limiting the number of periods on one side of the cavity, the light in the cavity will couple to that single radiating mode. Because this is the only way light can exit the cavity, the light source in the cavity can be 100% efficient [2].

For integrated photonics, the use of photonic crystals is very attractive. Because in a photonic crystal line defect the light is bound to the line defect, it can only travel back and forth along this defect for wavelengths in the photonic bandgap [4]. This creates a very effective waveguide. If we now bend this waveguide, as illustrated in figure 2.20, light propagating through it cannot leave the defect, but only continue to propagate through the bend or return along the waveguide [4, 58]. A

[3]The reflection from a three-dimensional photonic crystal is not necessarily specular. In most cases the reflected light is reflected (of diffracted) into a different directions, for the different orders of diffraction

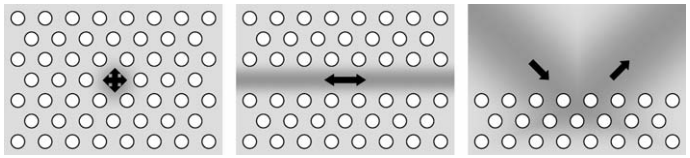

Figure 2.19: Defects in a 2-D photonic crystal. Left to right: a cavity point-defect, a line defect of smaller holes, a surface defect (which, for a 2-D periodic structure, is also a 1-D defect.

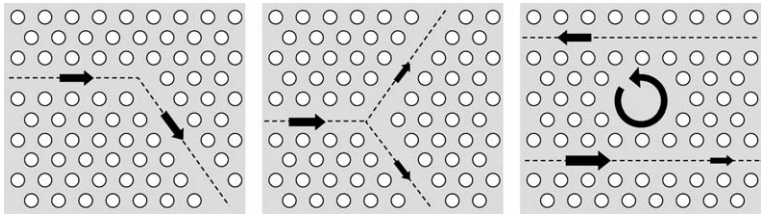

Figure 2.20: 2-D photonic crystal components. Left to right: a 60° bend, a splitter, a cavity drop filter.

good design of the bend geometry can maximise the bend transmission and eliminate the backreflection.

The same is true for the splitter in figure 2.20. Here the incoming light can be split in two equal parts, but here too, careful design is needed to eliminate unwanted reflection.

A last example in figure 2.20 is 2 waveguides with a cavity in between. Because the cavity is only slightly separated from the waveguide, light coming in from one waveguide can couple to the cavity, and if the wavelength *fits* the cavity, it will resonate and couple light to the other waveguide. Again, if properly designed, this could be done with very high efficiency, because the light cannot escape through the photonic crystal.

These examples show that defects in photonic crystals can indeed be very useful. However, we have also stressed the importance of good design. As we will show later in this chapter, designing photonic crystal components is nontrivial, requiring very intensive numerical calculations. For real-world applications, the modelling should be done in three dimensions, even for 2-D periodic structures.

2.4 Waveguides: Guided and Leaky Modes

Before bringing 2-D photonic crystals to the real world, we will take a closer look at some basic properties of waveguides, and how they are affected when we use them in combination with periodic structures.

Classical optical waveguides guide light by concentrating it into an area with a higher refractive index than its surroundings. This index difference is usually not high. In optical fibres, one of the most efficient type of optical waveguides, the index contrast is typically less than 1%.

As an example, we take the waveguide in the left part of figure 2.21, consisting of a high-index core and a lower-index cladding on either side. When we plot the dispersion relation of the waveguide and the uniform cladding, we see that for low frequencies, the dispersion of the waveguide is close to that of the cladding[4]. This is because the wavelength of the light is long, and only a small fraction feels the waveguide core. As the frequency increases, the dispersion curve of the waveguide moves away from the cladding line, and for high frequencies, most of the line is in the core, so the dispersion relation is similar to that of a uniform medium with the refractive index of the waveguide core. This is the ground mode, or zeroth-order mode of the waveguide.

All this time, the waveguide dispersion curve has not entered the area above the line defined by the in-plane dispersion relation of the cladding. As only the in-plane components of $\mathbf{k}$ are taken into account when we draw the cladding modes in the waveguide's dispersion diagram, we get a continuum of radiative modes in the band diagram with a lower boundary of radiative modes parallel to the slab waveguide. This line, formed by the dispersion relation of the in-plane radiation modes, is called the *cladding line* or the *light line*. Above this line, a continuum of modes exists that can freely propagate in the cladding. As long as the waveguide dispersion relation is outside of this continuum, the waveguide mode is truly guided by the refractive index contrast. At higher frequencies, we can see other modes appear.

If we now add a periodic corrugation to this waveguide, like in the middle part of figure 2.21, the waveguide mode becomes a Bloch mode, the dispersion relation becomes periodic, and can be folded into the first Brillouin zone, with the creation of a bandgap. The cladding is a uniform medium. Therefore, the waveguide mode that was once index-

[4]The wave vector $\mathbf{k}$ is now a wave number k. In the context of conventional waveguides the wave number is often indicated with the symbol β. However, to avoid confusion, we will continue to use k for the wave number throughout this work.

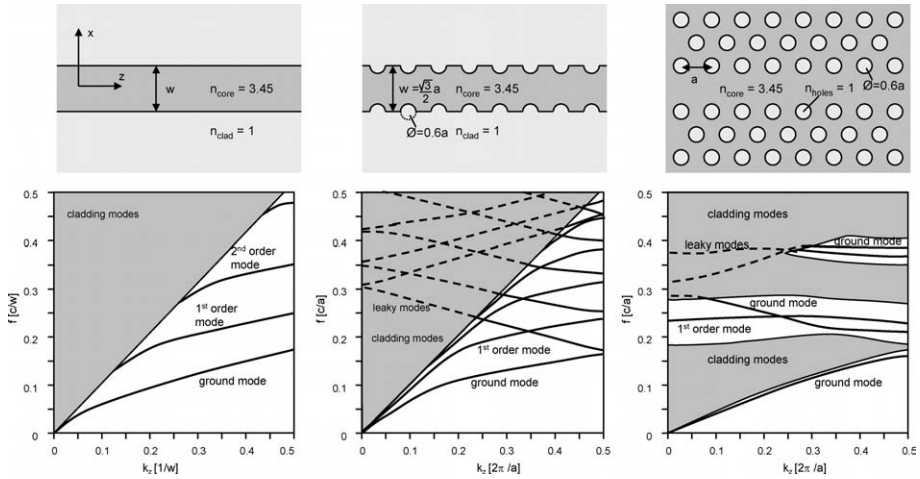

Figure 2.21: Waveguides and their dispersion relation for the TE polarisation. Left to right: Slab waveguide, Periodically corrugated slab waveguide, 2-D photonic crystal waveguide

guided over the whole spectrum is suddenly folded back above the light line of the cladding. At that point, the mode is not index guided anymore, and is coupled to the radiation modes by the periodic modulation.

If we now replace the cladding by a photonic crystal, we change the dispersion relation of the cladding as well, and create a photonic bandgap, as we have shown in the previous chapter. Now, we see that some of the folded-back bands of the ground mode and even one of the higher-order modes are now outside of the cladding any more because is has a photonic bandgap. Suddenly, the once *index-guided* mode has become a *PBG-guided* mode. When we plot the dispersion relation of the waveguide as a 1-D periodic structure, the propagation constants of the cladding modes of the 2-D periodic lattice are projected onto the k_z axis, creating a continuum of modes in the band diagram of the 1-D periodic waveguide.

2.5 Photonic Crystal Slabs

While 3-D photonic crystals show the most promise with respect to confinement, 2-D photonic crystals are much easier to fabricate. However, in a 2-D structure which is uniform in the third direction, light will diffract and radiate away. Therefore, light has to be bound in the plane

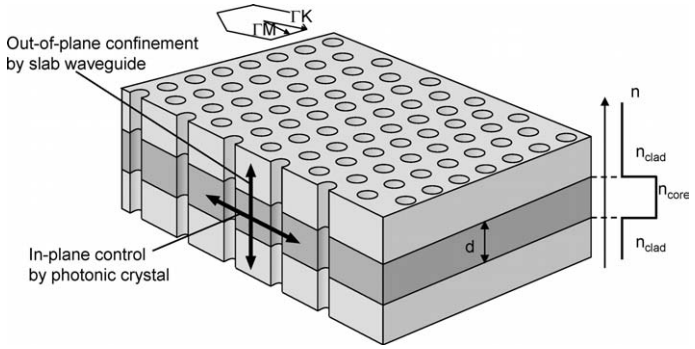

Figure 2.22: 2-D photonic crystal slab. In-plane, light is controlled by the photonic crystal, while in the vertical direction, it is confined by the layer with the higher refractive index.

of the photonic crystal. This can be done by combining the 2-D photonic crystal with a slab waveguide in the third direction, as illustrated in figure 2.22. Parallel to the layer structure of the slab waveguide, light is controlled by the 2-D photonic crystal, either propagating through a lattice or along a defect. In the vertical direction, light is bound to the slab waveguide by the layer with the higher refractive index, the slab waveguide *core*. We call these structures *photonic crystal slabs* or *planar photonic crystals*.

These 2-D photonic crystal slab waveguides are far more easy to make than the 3-D periodic structures. The layered slab waveguide can be made by growing or depositing a layer structure, and the photonic crystal can be made by defining the structure with lithography and etching them into the slab waveguide.

2.5.1 Materials Systems: GaAs, SOI, membranes

For a full bandgap, the photonic crystal should be made with a sufficiently high index contrast. For optical or telecom wavelengths, the best-suited materials are semiconductors, either from the III-V group like *Gallium Arsenide* ($GaAs$) and *Indium Phosphide* (InP) [5, 6, 59], or *Silicon* [7, 8].

The choice of material for the slab waveguide has a profound effect on the behaviour of a photonic crystal slab. If there is a strong index contrast between the *core* and *cladding* layers, the light will be tightly confined in the vertical direction. The extreme example of such high index contrast is a photonic crystal membrane, as illustrated in 2.23.

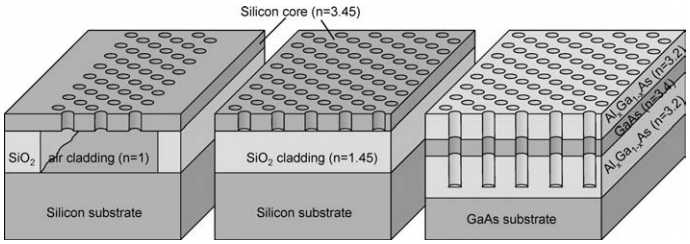

Figure 2.23: Types of 2-D photonic crystal slabs with various refractive index contrasts. Left to right: Silicon membrane, Silicon-on-insulator, $GaAs$-$AlGaAs$

This is made by etching holes in the core (mostly Silicon, but $GaAs$ is also possible), and then removing the cladding layer below. While this gives very good confinement in the vertical direction, the mechanical stability is rather poor, and because some structures cannot be made free-standing, the applications are limited.

A good alternative to semiconductor membranes is the combination of semiconductor and oxide. The most widely used example is *Silicon-on-insulator* (SOI). This material, which is also widely used for advanced CMOS applications, consists of a thin layer of single-crystal Silicon separated from the Silicon substrate by a layer of SiO_2, or glass. While the index contrast is less than that of membranes, the advantage is the structural stability and the compatibility with CMOS processing. An alternative to SOI is the use of a high-index $GaAs$ layer on top of a layer of $Al_xGa_{1-x}As$, which, by oxidation, can be turned into Aluminum Oxide ($AlOx$) with a lower refractive index.

While the high confinement in these structures bind the light tightly to the core, the behavior of these structures is somewhat different than that of the original perfect 2-D structures. Therefore, 3-D modelling is always required.

In the other case, when the index contrast in the vertical direction is low, the light will be spread out widely and resemble a plane wave. Because of this, the structure behaves more closely to the perfect 2-D structure, and approximate 2-D models are easier to use, by replacing the layered slab waveguide by a uniform material with an *effective index* calculated from the guided mode of the slab waveguide. The right part of figure 2.23 shows such a layer stack of $GaAs$-$Al_xGa_{1-x}As$, which has an index contrast of 3.4 to approximately 3.2, depending on the fraction x of Aluminum. Because the light is less confined vertically, the holes in

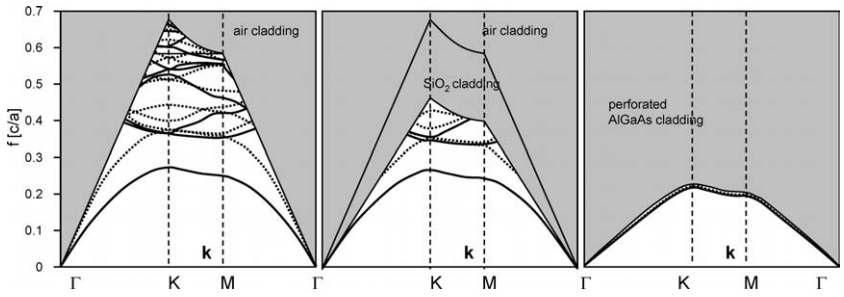

Figure 2.24: Band diagrams for the 2-D photonic crystal slabs from figure 2.23. Left to right: For a triangular lattice of holes in a membrane, in an SOI slab and in $GaAs$-$AlGaAs$. Quasi-TE modes are printed in a full line, while quasi-TM modes are dotted.

the structure should be etched deeply, so the slab mode feels the holes all over.

2.5.2 The Light Line in Photonic Crystal Slabs

Just like the simple examples in section 2.4, the photonic crystal slab is a waveguide in the vertical direction, where light should be confined to the core layer. Above and below the core is a semi-infinite cladding, which also supports its continuum of radiation modes. In the case of a cladding made of a uniform medium, like in free-standing membranes or SOI, the dispersion relation of the radiation modes is that of the uniform medium, with the magnitude of $\mathbf{k}$ linear with f or ω. In the case of a periodically modulated cladding, as in the right part of figure 2.23, this is no longer the case.

Figure 2.24 shows the band diagram for the structures of figure 2.23, and we have added the light line of the upper and/or lower cladding. Because in the membrane the refractive index of the cladding is much lower, the light-line is much steeper in the band diagram. As the index contrast lowers because of the higher cladding index, the light line drops.

We can see from the band diagrams that the photonic crystal has allowed modes below the light line. This is because of the slab waveguide, the index of the core is higher than the index of the cladding with its radiative modes. Also, we see that for the membrane and the SOI structure, a significant portion of the photonic bandgap is below the light line. However, as the index contrasts decreases, this region becomes smaller and smaller.

This region is of great importance to create useful components in 2-D photonic crystal slabs. When there is a photonic crystal mode, or *Bloch mode* below the light line, it cannot not couple to the radiation continuum unless the periodicity of the photonic crystal is broken [39, 40]. On the other hand, a Bloch mode above the light line will always have a small radiative component. In that sense, it behaves more like a resonance than a truly guided mode. In the next chapter, where we will examine the loss mechanisms in photonic crystal slabs, this will be explained in more detail.

2.6 Photonic Crystal Slab Waveguides

Just like in real 2-D photonic crystals, we can add defects to 2-D photonic crystal slabs. By removing or changing one or more rows of holes, we create a line defect along which light can propagate. This guided mode is now confined in-plane by the photonic crystal, and out-of-plane by the slab waveguide.

A straight line defect in a 2-D photonic crystal is in effect a 1-D periodically structured waveguide. Therefore, the dispersion relation will will also have frequency ranges where transmission is allowed and ranges where transmission is prohibited. Figure 2.25 shows the band diagram of a photonic crystal slab waveguide in Silicon-on-insulator, overlaid on the band diagram of the photonic crystal slab without defect. The waveguide has a number of guided modes with different properties. The lowest branch of the dispersion curve indicates the ground mode of the waveguide. Because the waveguide consists of a removed row of holes, the average index is higher than that of the surrounding photonic crystal. So, even though the waveguide mode is not in the photonic bandgap of the photonic crystal, it is guided because of its higher refractive index. Therefore, it is an *index-guided mode*, in both the vertical direction as the horizontal direction. In a sharp bend, however, it is not guaranteed that this index-guided mode will be bound to the line defect, because there is no PBG to prohibit the mode from propagating through the photonic crystal. On the other hand, the mode is below the light line, and is therefore a truly guided Bloch mode that will not couple to the radiation continuum.

We also observe a higher branch of the waveguide dispersion relation. This branch is quite horizontal and is *PBG-guided* in-plane. Because it is also below the light line, it is also index-guided in the vertical direction, and therefore a truly guided Bloch mode. It has, however, a

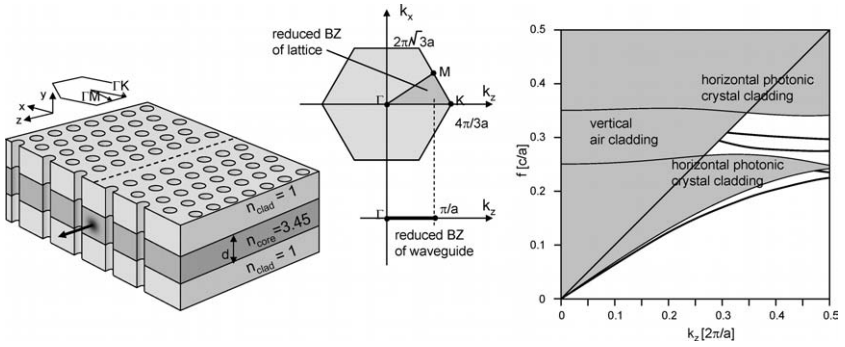

Figure 2.25: A 2-D photonic crystal slab waveguide in a Silicon membrane ($n_{core} = 3.45$ and $n_{clad} = 1$) and its band diagram. The structure has a mode that is index-guided in the vertical and horizontal direction, and two modes that are PBG-guided in the horizontal direction but index-guided in the vertical direction

very flat dispersion relation, meaning that the mode will travel slowly through the waveguide (i.e. it will bounce back and forth a lot.) This also makes the mode very narrow-band, and only useful over a limited wavelength range. At a certain point, the mode crosses the light line, and from that point on it ceases to be truly guided, and it will couple to the radiation modes.

In order to make full use of photonic crystals for waveguides, a number of conditions should be fulfilled. First, in order to avoid coupling to radiation modes, the guided mode should be below the light line. Secondly, it should preferably be PBG-guided in order to make use of the PBG-based confinement in bends, splitters and cavities. These two conditions are not straightforward to combine, as we will discuss in detail in chapter 3.

2.7 Polarisation

We have seen that there are different conditions for a photonic bandgap for either TE or TM. In a photonic crystal slab, the fabrication requirements make it nigh impossible to create a structure which has a PBG for both polarisations, so a choice has to be made.

In this work we will design nanophotonic components for the TE-polarisation. This choice is motivated by a number of factors:

- For the TE polarisation, the optimal structure for a large PBG is an interconnected network of high-index veins. A triangular lattice of air holes in a high-index slab waveguide is prime example of this structure. This interconnectedness of the slab core layer is beneficial for the vertical confinement.
- In chapter 4 we will see how the TM mode is more susceptible to losses, both from roughness on the sidewalls and from leakage to the substrate.
- In the next chapter, we will discuss a vertical fibre coupling structure. This structure only works for the TE polarisation.

Because of this choice, we will have to make sure that the TM polarisation is not excited in our structure. In the symmetric slab structure of figure 2.24, the TE and TM polarisations do not interact. However, if the vertical symmetry is broken, as is the case in Silicon-on-insulator, the TE and TM modes can couple. Therefore, we have to be careful in our designs that this coupling is suppressed where possible.

2.8 Modelling Tools

For the design of photonic crystal devices, good modelling tools are needed that efficiently solve Maxwell's equations. Because photonic crystals consist of wavelength-scale features with a very high refractive index contrast, it is not possible to apply approximations that are used in conventional, large photonic structures. Maxwell's equations should be solved as rigourously as possible, and in their full vectorial form.

There are a number of tools that offer this capability, of which we will briefly discuss three. *Finite Difference Time Domain* or *FDTD* is the most widely used simulation technique for photonic crystals [9]. It relies on a brute-force discretisation of Maxwell's equations. However, it is not very efficient, and inadequate for some problems. On the other hand, *plane-wave expansion* (PWE) and *eigenmode expansion* (EME) try to describe the electromagnetic field as a superposition of simple wave functions, like plane waves or the eigenmodes of the structure [13].

2.8.1 FDTD

The finite difference time domain (FDTD) method replaces the partial differentials in Maxwell's equations by finite differences, hence its

name [9]. This straightforward approach makes it possible to simulate a wide variety of structures. Basically, a structure is defined and discretised on a spatial grid. A time-dependent source is then placed somewhere in the structure, and the **E** and **H** fields are iteratively calculated throughout the structure for each subsequent time step. These fields can then be recorded in certain places for later analysis.

Even though this scheme can be highly optimised, the method is still very resource intensive, requiring huge amounts of physical memory to store the fields throughout the structure. Also, the intricacies of the method dictate that the time step should be linearly dependent on the spatial discretisation, so the simulation of a finely meshed structure requires a large amount of time steps, and thus a large calculation time. On the other hand, FDTD's efficiency is independent from the structure geometry, making it an ideal tool for the modelling of arbitrary structures. Since its introduction, FDTD has been optimised for various purposes, including the calculation of band diagrams [60].

Because FDTD can only simulate structures of a limited size, it is not straightforward to model open structures that can radiate light. Therefore, good boundary conditions are needed to absorb the light emitted by the simulated structures and eliminate artificial, and therefore unwanted reflection from the boundaries of the simulation space. *Perfectly matched layers* (PML) offer this possibility, by 'coating' the hard boundaries with an absorbing layer that doesn't reflect the incoming radiation [10].

2.8.2 Plane-wave Expansion

Instead of describing the electromagnetic field by storing its value in each point, one can also describe it as a superposition of basis functions. Plane waves, the natural solution for Maxwell's equation in a uniform medium, form such a set of orthogonal functions, and are very simple in use [61]. When used in the frequency domain (i.e. for monochromatic light with frequency f), the Maxwell equations can be reduced to a rather simple eigenproblem. The resulting eigenmodes then come out as a superposition of plane waves, resulting in the localised modes of the structure. The accuracy of the method is determined by the number of basis functions, in this case plane waves, are used in the calculations.

Just as with FDTD, the simulation space of this technique is limited, but the boundary conditions are periodic, so the structure is repeated in all simulation directions. While this limits the applicability, it makes

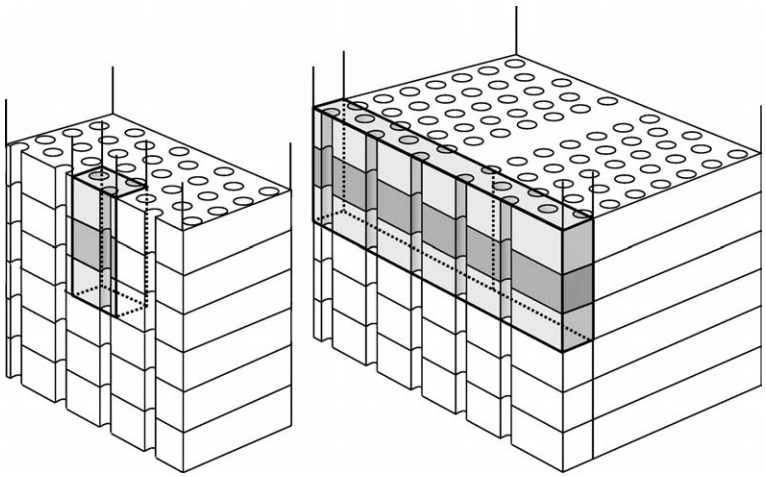

Figure 2.26: Simulation cells for plane wave expansion. Left: A photonic crystal slab. Right: a photonic crystal slab waveguide, which requires a large supercell.

the technique well-suited for photonic crystals. Figure 2.26 shows how the simulation space can be chosen to encompass one photonic crystal period. Plane-wave expansion is the most widely used method for the calculation of band diagrams of photonic crystals. For photonic crystals with defects, a supercell is used, encompassing many periods of the photonic crystal, as illustrated in figure 2.26. However, this requires more intensive calculations, as more plane-waves are needed to describe the electromagnetic field accurately. The supercell should also be sufficiently large so that the structures cannot couple to their periodically repeated neighbours.

Because of the requirement of periodic boundary conditions, it is also difficult to simulate open structures. However, one can also introduce perfectly matched layers within the structure, effectively absorbing the light. Even though the boundary conditions remain periodic, the structures are effectively isolated from each other by PML [62, 12]. While in the right part of figure 2.26 the out-of-plane radiation can couple from the photonic crystal slab to its neighbours, the PML would then absorb this radiation.

An example of a plane-wave expansion tool is the MPB package developed at MIT [11]. It uses plane waves to calculate the Bloch modes for periodic structures in 1, 2 or 3 dimensions, and we used it for the band diagrams in this chapter. Because we can deduce much informa-

tion from this even without knowing the actual fields in all the points, plane-wave expansions is in many cases more effective than FDTD.

Plane waves form the most simple set of orthogonal basis functions. However, they are the eigenmodes of a uniform medium, and therefore continuous functions. Because it is not easy to satisfy the (often discontinuous) boundary conditions on the interfaces between materials, they are not ideally suited to describe the electromagnetic field in high-contrast structures like photonic crystals. In order to get an accurate description, many plane waves should be taken into account. This reduces the effectiveness of the technique.

2.8.3 Eigenmode Expansion

Instead of using plane waves, one can choose a better set of basis functions. If a structure is invariant along the propagation axis, like many ridge waveguides, the eigenmodes of the cross section can be used as a set of basis functions. These eigenmodes behave independently of each other in that particular section, each one propagating with its own propagation constant k.

When the structure changes in the propagation direction, it can be discretised in piecewise constant cross sections. In each section i, the field is described in the eigenmodes of that section Ψ_{ij}, and the propagation of each eigenmode is calculated independently of the others. At the interfaces between structures, a scattering matrix is calculated describing the coupling between the eigenmodes of both sections. This way, the electromagnetic field can be decomposed into the eigenmodes of section $i + 1$.

Once the eigenmodes of a particular section are calculated, they can be reused for other sections with the same profile. This is also the case for recurring interfaces between sections, where the scattering matrices are identical. This makes eigenmode expansion very suitable for structures that are periodic in the propagation direction, like photonic crystal waveguides.

Just as with the plane-wave expansion and FDTD, the PML boundary conditions can be used to eliminate reflections from the boundaries of the simulation space. While with FDTD PML is often implemented as an anisotropic material, eigenmode expansion treats this absorbing boundary condition as a layer with complex thickness [13].

Because the eigenmodes are much better fit for the simulated structure, only a relatively small number of modes is needed. However,

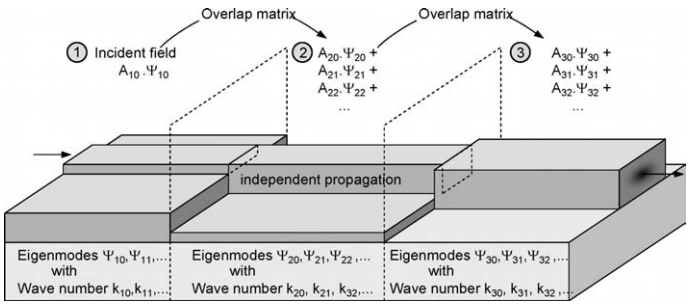

Figure 2.27: The principle of eigenmode expansion: In each section i, light is described in the eigenmodes Ψ_{ij} of that section that propagate independently with their own propagation constants k_{ij}. At the interface between sections, a scattering matrix is calculated from the overlaps between the eigenmodes of both sections to map the fields onto the eigenmodes of the new section.

in contrast with plane waves, the modes should be calculated first, which is not trivial when dealing with absorbing materials. This is also the case when using PML. Because of this, eigenmode expansion does not automatically guarantee a correct solution, though the evolution of mode solving techniques make the method more and more robust.

We have used both commercially available eigenmode expansion tools, like *Fimmwave-Fimmprop3D* from Photon Design[5], as the open-source package *CAMFR*[6], developed by the *photonics* group at the INTEC department of Ghent University [13, 63]. Both software packages have their strengths and weaknesses, especially when high-contrast photonic crystals are involved. While Fimmwave has a graphical user interface, CAMFR is easily scriptable, and has the capability to calculate band diagrams of 1-D periodic structures. Because of this combination of both efficiency and versatility, it is the most-used technique for the simulations in the following chapters.

2.9 Conclusion

In this chapter, we have explained where photonic crystals get their peculiar properties from. Because of the periodicity, they have frequency or wavelength ranges known as *photonic bandgaps* where no propagation of light is allowed in the structure. Generally, the higher the index

[5]Photon Design web site: *http://www.photond.com*
[6]CAMFR website: *http://camfr.sourceforge.net*

contrast of the structure, the larger this bandgap becomes. 1-D photonic crystals are most often used as frequency selective mirrors. 2-D and 3-D photonic crystals, which only have an omnidirectional photonic bandgap for sufficiently high index contrast and certain symmetries, become very useful when adding defects. These defects create localised modes in the structure. Line defects can be used as waveguides, while point defects are ideally suited for high-quality cavities.

2-D photonic crystals are the easiest to fabricate, especially when combined with a slab waveguide in the third direction. However, it is not trivial to create conditions for a good waveguiding device, as coupling to the radiation modes in the slab waveguide's cladding can introduce severe losses. In the next chapter, we will look deeper into these loss mechanisms.

Because of their high refractive index contrast and wavelength-scale features, photonic crystals are exceptionally difficult to model. While FDTD provides the most versatile method, it is very numerically intensive. We have therefore chosen to use plane-wave expansion and eigenmode expansion, depending on the application. These techniques are numerically much more efficient.

Chapter 3

Nanophotonic Components

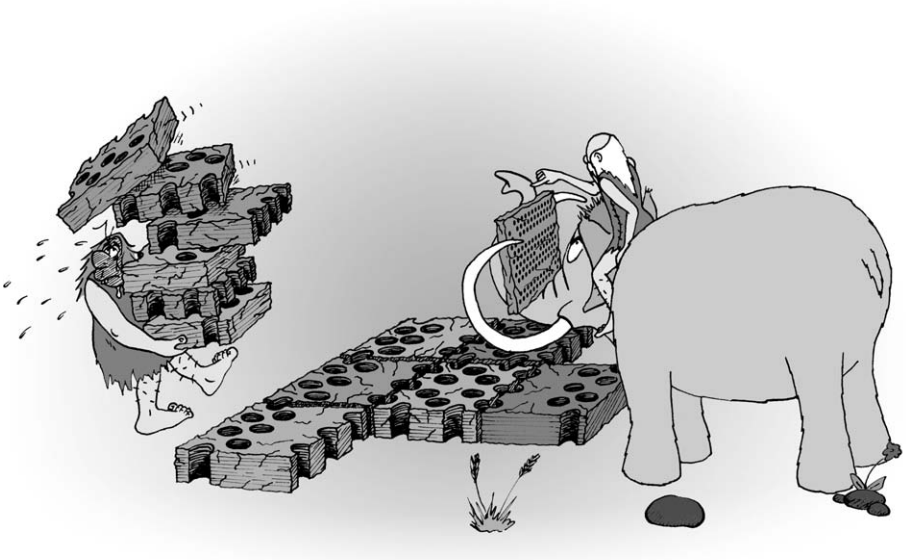

"What do you mean: it's really simple?"

Chapter 3

Nanophotonic Components

In nanophotonic integrated circuits, waveguides are an extremely crucial component. Not only are they necessary for interconnects, but many functional elements are also based on waveguides. Therefore, good waveguides are a prerequisite for further integration. Nanophotonic waveguides need to confine light into a small area. In order to do this and to allow for sharp bends, high index contrast is required.

However, nanophotonic waveguides don't interface easily to the outside world. Single-mode fibre has a large core, small refractive index contrast and a small numerical aperture. In this it represents everything a nanophotonic waveguide is not. Coupling light from one to the other is therefore a nontrivial challenge.

In this chapter, we will discuss a number of waveguide designs, as well as a number of components based on nanophotonic waveguides, such as resonators. We will also discuss some components with a relatively small footprint that allow coupling of nanophotonic waveguides to a single-mode fibre. Many of these designs are adaptations of designs published in literature or the result of collaborations with other groups.

3.1 Material Choice: Silicon-on-Insulator

As we have discussed in chapter 2, we can use a variety of materials for photonic crystal slabs. For most of this work, we opted for *Silicon-on-insulator* (SOI). The main reasons for this choice are its very good optical properties at telecommunications wavelengths and its compat-

ibility with CMOS fabrication processes, discussed in chapter 5. The fabrication technique for SOI wafers is also discussed in chapter 5.

3.1.1 Optical Properties

SOI consists of a thin Silicon layer on top of an oxide cladding layer deposited on a bare Silicon wafer. Silicon-on-insulator is a very good material for optical waveguides. With its Silicon core $n = 3.45$ and its oxide cladding ($n = 1.45$) it has a high vertical refractive index contrast. Also, both the Silicon and the oxide are transparent at the telecom wavelengths of $1.3\mu m$ and $1.55\mu m$.

As will be discussed in chapters 4 and 6, we chose an oxide thickness of $1\mu m$ to reduce substrate leakage. The thickness of the core was chosen to be $220nm$, in order to keep the slab waveguide single mode for the TE polarisation.

3.1.2 Deep or Shallow Etching

When we want to fabricate nanophotonic components in Silicon-on-insulator, we can choose between two etch procedures, each with its merits and drawbacks. Figure 3.1 illustrates both procedures for a photonic crystal slab. When we etch both the top Silicon and the underlying oxide, we increase the refractive index contrast even more, because we replace the oxide ($n = 1.45$) with air ($n = 1$). Also, the bottom of the photonic crystal holes is far away from the guided mode, which reduces a possible source of scattering. An additional advantage is that the vertical layer structure becomes more symmetric, reducing the interaction between TE and TM modes. However, as we will illustrate in the chapter 5, this approach introduces significantly more sidewall roughness, which, we will see in chapter 4 is a principal source of losses.

When listing all these advantages of deep etching, is there any reason to etch only the Silicon, as illustrated in the right part of figure 3.1? When we etch only the Silicon, the sidewall roughness due to etching can be significantly reduced. Also, in order to eliminate the interaction between TE and TM, oxide can be deposited on top of the Silicon core after etching. This deposition, briefly discussed in chapter 5, works better with shallow holes than with deep holes.

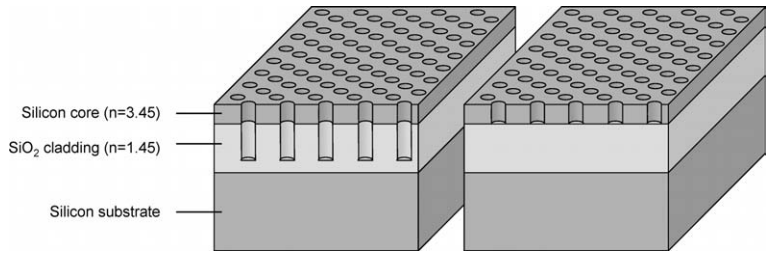

Figure 3.1: SOI Photonic crystal slabs with deep and shallow etch.

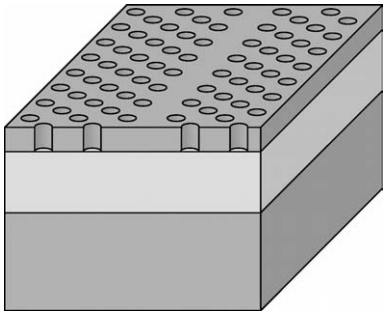

Figure 3.2: W1 SOI Photonic crystal slab waveguide

3.2 Photonic Crystal Waveguides

When the concept of photonic crystal waveguides was introduced, the design of a waveguide seemed as simple as removing a row of holes. This type of structure is illustrated in figure 3.2 for Silicon-on-insulator. As we will see in the next sections, this design has indeed a guided waveguide mode, but its properties are ill suited for pure waveguiding. Therefore, more optimised designs are needed.

Apart from pure waveguiding, photonic crystal slabs can also be used for other purposes. One significant aspect of periodic structures is the appearance of wavelength regions with a very flat dispersion curve, and therefore a low group velocity, especially near the band edges [64]. In these regions, light is coupled strongly between the backward and forward propagating direction, and travels slowly through the waveguide. If the waveguide has active properties, like gain or nonlinear effects, the interaction time of the light with the material is significantly enhanced for these *slow waves*.

3.2.1 Requirements for Photonic Crystal Waveguides

If we want to use photonic crystal slabs for simple waveguiding, our designs should fulfill a number of conditions, illustrated in figure 3.3:

- As discussed in the previous chapter, we want to use the TE polarisation. In SOI, the coupling to TM is always present, but if possible, we want to minimise this.
- We want to have a **PBG guided mode**, so we can make use of the photonic bandgap to prohibit radiation in bends, splitters, and other defects. Therefore, the lattice photonic bandgap should be sufficiently large. This condition, together with the previous condition, suggests the use of a triangular lattice of holes.
- To reduce losses due to out-of-plane scattering (see chapter 4), the mode should operate below the light line of both the air and oxide cladding.
- Preferably, we want a single mode waveguide. This means that either we work below cut-off for the first-order mode, or the first-order mode has a stop-band at the operating wavelength.
- We want to use the waveguide for telecommunication, and therefore it should be able to transport a number of wavelength channels. The guided mode has to have sufficient bandwidth for this purpose, i.e. at least 50nm.
- For waveguiding purposes, we want to avoid strong dispersion and slow waves. The dispersion relation of the guided mode should have a sufficiently steep slope, which remains constant over the required wavelength range, in order to keep a constant, high group velocity.
- In order to integrate the waveguide with other photonic crystal based components, the lattice of the cladding should be kept in place. This means that we cannot shift parts of the lattice around, as that would make it difficult to design bends and splitters, or align different waveguides with respect to each other.

Some of these conditions contradict. For instance, in order to get a large PBG, a strong index contrast is needed between the SOI stack and the etched holes. This index contrast is higher with a thicker core

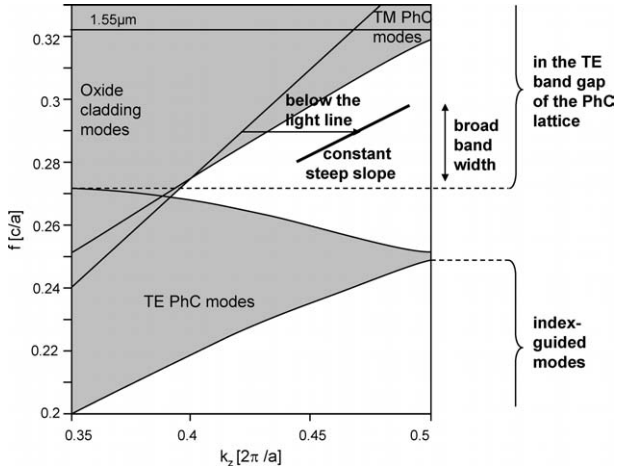

Figure 3.3: Requirements for an SOI photonic crystal waveguide mode: Below the light line, in the PBG of the photonic crystal, a constant, steep slope, and sufficient bandwidth.

layer. However, in a thicker core layer it is more difficult to create a single-mode waveguide.

The lateral effective index contrast between waveguide and etched holes could also be increased by using a higher cladding index, instead of oxide. This is the case in conventional III-V heterostructures like $GaAs - AlGaAs$ layer stacks. But as we have seen on page 40, this in turn will lower the light line, reducing the useful wavelength range.

In a conventional so-called *W1* waveguide, where we simply remove one row of holes, these conditions are not met. In this waveguide, the average refractive index of the waveguide is higher than the surrounding lattice. Therefore, this is an *index-raised* defect, where the lowest branch of the ground mode is not in the photonic bandgap of the lattice. Only the higher branches are in the PBG, and these branches have a small slope, making them unsuitable as a waveguide. Therefore, better designs are needed.

For our designs, we chose a lattice pitch of $500nm$ and an originally a hole diameter of $300nm$, and later $320nm$. While the useful area below the light line is larger for smaller pitches (between $400nm$ and $460nm$, these structures are harder to fabricate reliably.

3.2.2 Impact of Fabrication Techniques

Additional design limitations are imposed by the fabrication. For example, it is very difficult to define lattices with very large holes that almost touch each other. In this situation, the smallest feature, i.e. the narrow stripe of material in between the holes, becomes too small to resolve with most pattern writing techniques.

For the devices fabricated in the work, we made use of deep UV lithography. A detailed discussion of this technique with its advantages and drawbacks is given in chapter 5. However, we can already mention that the resolution of this technique is strongly limited by the pitch of periodic structures. From a lattice period of $400nm$ down, it becomes increasingly difficult to fabricate structures correctly and reproducibly. Therefore, we based our structures on a lattice with $500nm$ period. To get a sufficient PBG around $1.55\mu m$, a hole diameter between $300nm$ and $340nm$ is needed, depending on the choice of the core thickness d_{core}. In our device wafers, we chose d_{core} to be $220nm$.

As we can see in figure 3.3, these parameters are far from optimal from a design point-of-view. Because the pitch is rather large, the operating wavelength of $1.55\mu m$ is located above the most interesting region. In fact, in the asymmetric SOI structure, there is a possible coupling between TE and TM modes in the continuum of TM lattice modes.

Also, as we have already mentioned on page 55, we have the choice between etching both silicon and oxide or etching the oxide. Figure 3.4 shows the band diagrams for both structures, illustrated in figure 3.1. The core thickness is $220nm$, the lattice pitch is $500nm$ and the hole diameter is $300nm$. We can clearly see that, by etching the underlying oxide, the light line of the bottom cladding is raised. Also, because the bottom cladding is a periodic structure, the light line ceases to be a straight line. The steeper light line creates additional freedom for waveguide designs by freeing more useful 'area' in the PBG of the lattice.

From this, it is obvious that deep etching has considerable advantage with respect to design. However, as will become clear in chapter 5, etching the oxide poses a considerable challenge from a fabrication point-of-view.

3.2.3 Waveguide Designs

In this section, we will discuss a number of photonic crystal slab waveguide designs in SOI. We assumed the deeply etched SOI structure

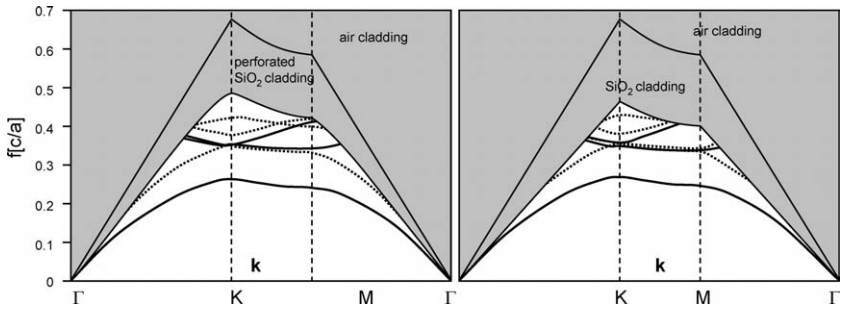

Figure 3.4: Band diagrams for the SOI Photonic crystal slabs from figure 3.1 on page 55, with a lattice pitch of $500nm$ and a hole diameter of $300nm$.

for its better properties. However, most designs are also valid for the Silicon-only etched structure.

Most of these waveguide designs are adaptations from designs published in literature, where we adapted the parameters for our layer structure [65, 7].

Simple W1 Waveguide

We have already briefly discussed the simple W1 waveguide, consisting of a single row of holes removed from the lattice. This structure, illustrated in figure 3.5, seems a simple way of creating a good waveguide. However, when we look at the associated band diagram, we see that this waveguide does not satisfy a number of our waveguide criteria:

- For the TE-like modes, the part of the dispersion curve below the light line is very flat. This makes the waveguide narrow-band.
- In the PBG, the waveguide supports multiple modes, some of which are above the light line and therefore lossy.
- The interaction with the TM ground mode gives rise to a mini-stop-band, i.e a narrow wavelength region where neither one nor the other is guided. This is clarified further at the end of chapter 4.

In order to remedy these drawbacks, changes need to be applied to the concept of the W1 waveguide.

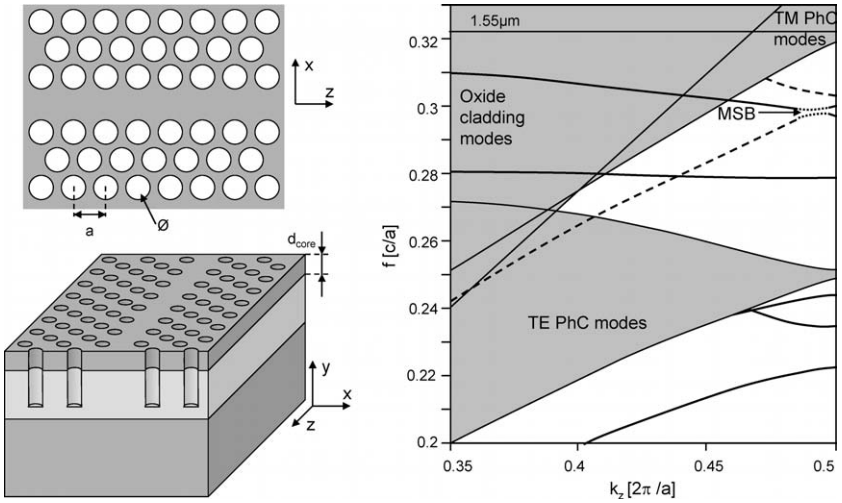

Figure 3.5: W1 photonic crystal waveguide and band diagram. TE-like modes are full line, TM-like modes are dashed. Mixed modes are dotted. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$, SOI core thickness $d_{core} = 220nm$.

W3 Waveguide

When we remove 3 rows of holes, we create a W3 waveguide. Because the optical volume of the defect is even larger than that of the W1 waveguide, the W3 waveguide is heavily multi-mode, with a large number of modes above the light line. Also, the ground TE mode will not only interact with TM modes, but also with higher-order TE modes, creating a so-called mini-stop-band. Because this MSB couples light between 2 guided modes, it can be used as a wavelength-selective element, coupling light at a single wavelength from the ground mode to the higher-order mode, and leaving the rest of the light unchanged. However, as a simple waveguide, its multimode nature makes the W3 unattractive.

Reduced-Core-Volume Waveguides

In order to make the W1 waveguide single-mode, we have to reduce the optical volume of the waveguide core. This can be done by reducing the core dimensions. The W0.7 waveguide, proposed by Notomi et al. [66] at NTT, does this by shifting the lattice of the waveguide so the distance between the hole centers on each side of the waveguide is 0.7 times the distance in the original W1 waveguide, which is $\sqrt{3}a$. This structure,

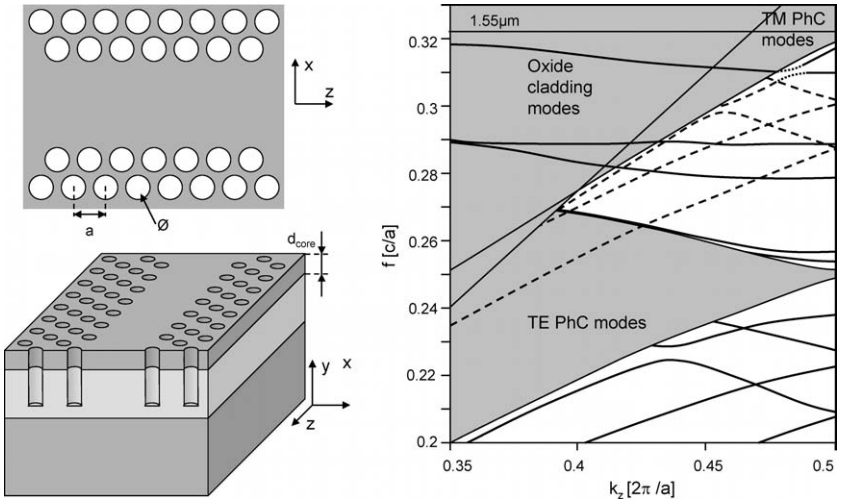

Figure 3.6: W3 photonic crystal waveguide and band diagram. TE-like modes are full line, TM-like modes are dashed. Mixed modes are dotted. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$, SOI core thickness $d_{core} = 220nm$.

illustrated in figure 3.7, supports a single guided mode in the photonic bandgap. With this structure, Notomi et al. demonstrated record-low propagation losses of $6dB/mm$ in 2002 [14].

However, while this structure has very good waveguiding properties, it uses a shifted lattice to reduce the core volume. Because of this lattice shift, it becomes difficult to integrate the waveguide into larger circuits with bends and splitters without introducing lattice discontinuities at some points. Therefore, alternative ways of reducing the core volume are desirable.

Changed-Hole W1 Waveguides

The optical core volume is determined by both the physical dimension of the core and the refractive index of the material. By removing part of the core material, the optical core volume can also be reduced. Instead of removing a hole, we can change the hole diameter of a row of holes.

To do this, we can either enlarge or reduce the hole size. Figure 3.8 shows a W1 waveguide where the diameter of the core holes $\oslash_{def} = 200nm$ instead. Because the introduction of holes in the W1, the average refractive index of the waveguide core decreases, and therefore

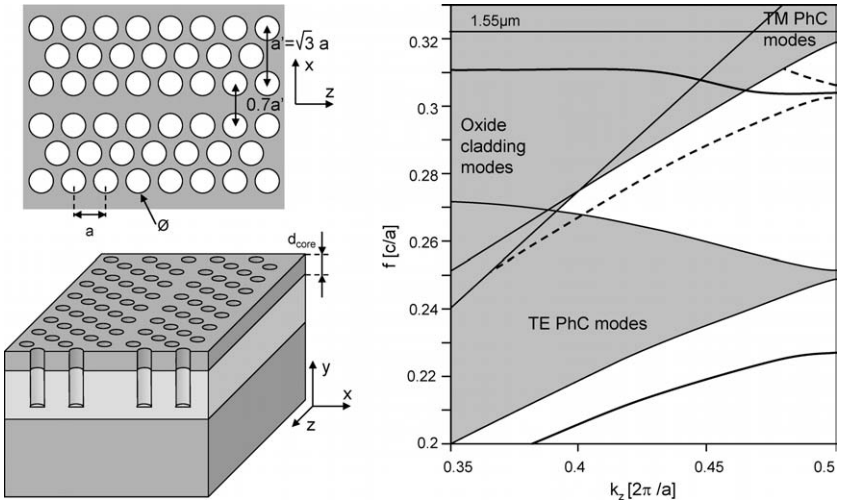

Figure 3.7: W0.7 photonic crystal waveguide and band diagram. TE-like modes are full line, TM-like modes are dashed. Mixed modes are dotted. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$, SOI core thickness $d_{core} = 220nm$.

the bands are pushed upward compared to those of a simple W1 waveguide.

Although the waveguide has a guided mode with an acceptable slope in the PBG below the light line, this is the first order mode. Therefore, this mode is antisymmetric in the x direction, and it is not straightforward to couple light into such a mode.

The folded branch of the ground TE mode, which in the simple W1 interacted with the ground TM mode, is now pushed upward into the TM lattice modes, where it will interact with the TM modes because of the vertical asymmetry of the SOI layer stack.

The smaller holes create a waveguide which has still a refractive index that is higher than the surrounding lattice. Therefore, the waveguide always supports an index-guided mode. If we create a row with larger holes than the surrounding lattice, we have an *index-lowered* defect, and index-guiding is no longer possible. In figure 3.9 we can see in the band diagram that there are no more index-guided modes in the region below the TE lattice modes. We also see that there is a folded branch of the ground mode with reasonably large bandwidth.

While this structure has good qualities for a waveguide, it also has some drawbacks. First, the structure has a very large amount of sidewall surface right in the core area where most of the light is confined.

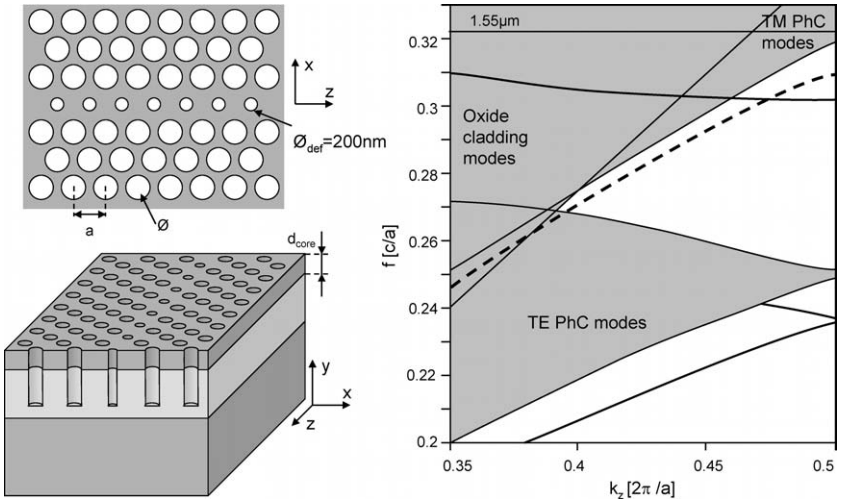

Figure 3.8: W1 photonic crystal waveguide with small defect holes and band diagram. TE-like modes are full line, TM-like modes are dashed. Mixed modes are dotted. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$, SOI core thickness $d_{core} = 220nm$. The defect holes are smaller than the lattice holes with $\oslash_{def} = 200nm$.

This makes it extremely prone to scattering at sidewall roughness. Secondly, the larger holes imply a very narrow Silicon wall in between the holes. Such a small dimension is very difficult to fabricate reliably. Also, small variations in the wall thickness will have a serious impact on the dispersion relation, shifting the mode up and down the band diagram. As will briefly be discussed in chapter 4, this makes the waveguide vulnerable to losses due to band shifting.

Trench W1 Waveguides

An alternative way of creating an index-lowered defect without introducing very narrow Silicon lines is by replacing a row of holes by an etched trench. Although it seems that such a structure cannot confine light in the vertical direction, the band diagram in figure 3.10 shows the TE ground mode to be below the light line. However, the low index in the center of the waveguides gives rise to a somewhat deformed mode shape. This will make it difficult to effectively couple light into this waveguide.

Although this design eliminates the need for thin Silicon walls, it still creates a large amount of sidewall, making the waveguide less tol-

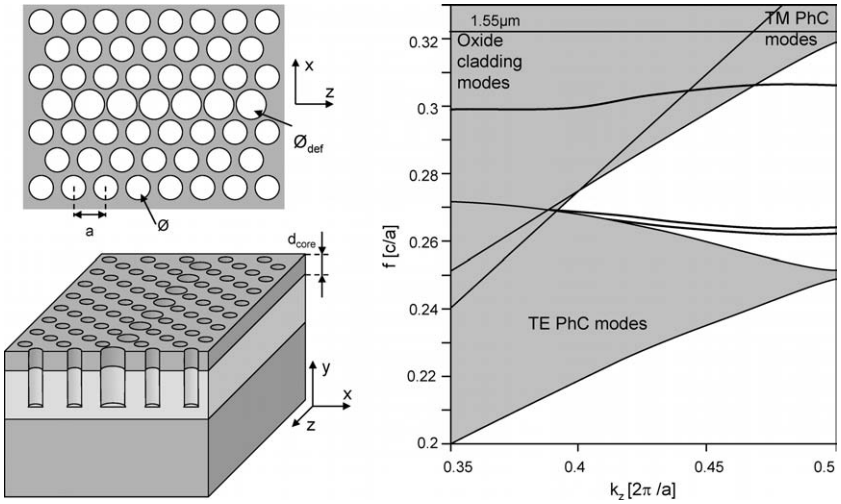

Figure 3.9: W1 photonic crystal waveguide with larger defect holes and band diagram. There are only guided modes for the TE polarisation. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$, SOI core thickness $d_{core} = 220nm$. The defect holes are larger than the lattice holes with $\oslash_{def} = 460nm$.

erant for sidewall roughness. Also, the optimal trench width of $120nm$ is rather small, which will make this waveguide also difficult to fabricate.

Photonic-Crystal-Embedded Wires

Instead of a single trench, one can also use a double trench to create the defect [67]. This structure, illustrated in figure 3.11, is actually a photonic wire embedded in a photonic crystal. This structure, proposed by Fan [15], combines the good waveguiding properties of a photonic wire with the in-plane PBG confinement of a photonic crystal slab.

W1 Waveguides with Changed Border Holes

An alternative way of reducing the optical core volume is by changing the border holes of a W1 waveguide. An example is shown in figure 3.12. Here, the holes near the defect are enlarged. While this widens the defect to three rows of holes instead of one, the actual waveguide core is smaller. However, this design also introduces additional sidewall surface. If we look at the band diagram, we see that it does have good

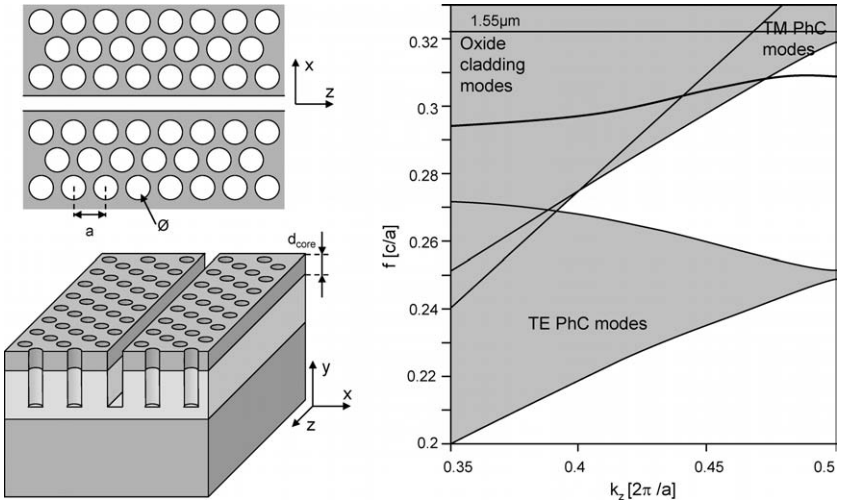

Figure 3.10: W1 photonic crystal waveguide with a defect consisting of a single trench. There are only guided modes for the TE polarisation. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$, SOI core thickness $d_{core} = 220nm$. The trench width $w_{def} = 120nm$.

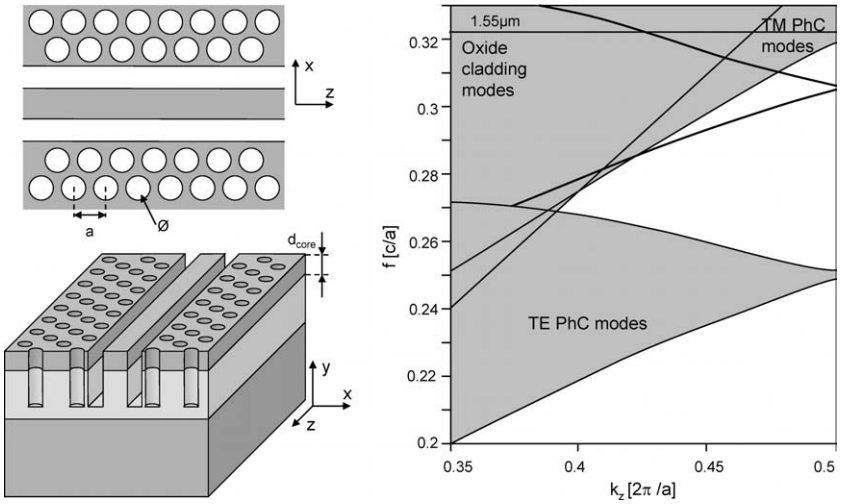

Figure 3.11: Embedded photonic wire waveguide and band diagram. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$, SOI core thickness $d_{core} = 220nm$. The trench width $w_{def} = 350nm$, and the wire width $w_{wire} = 330nm$

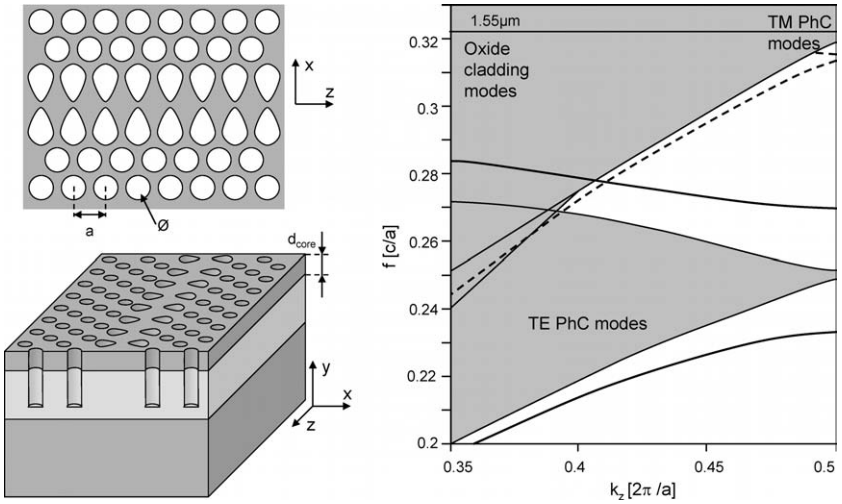

Figure 3.12: W1 photonic crystal waveguide with deformed border holes and band diagram. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$, SOI core thickness $d_{core} = 220nm$.

waveguide characteristics. The guided mode has a steep and uniform slope over a large bandwidth.

Conclusion: Waveguide Designs

While all these designs have their peculiar strengths, there is no design that convincingly satisfies our requirements for a good waveguide. While Notomi et al. have already demonstrated low propagation losses for the W0.7 waveguide [14], it has a shifted lattice, which makes it difficult to integrate many components with this waveguide. Other techniques to reduce the optical volume of the waveguide core invariably introduce additional sidewall surface, making the design very sensitive to fabrication quality.

3.2.4 Bends and Splitters

The advantage of photonic crystal slab waveguides is that the light is confined in-plane by the photonic bandgap. From a simple view, when we introduce a bend or splitter in the waveguide, the light cannot radiate into the cladding, but should follow the bend or split into two equal parts in the splitter [4]. However, there are two mechanisms that reduce the efficiency:

- The light is bound to the waveguide defect. At the bend or splitter, it can either continue its way through the bend or be reflected.
- In an SOI photonic crystal slab, the light is only perfectly bound to the waveguide as long as the translational symmetry of the waveguide is unbroken. At a bend or splitter, the light can couple to the radiative modes and be scattered into the vertical cladding.

Good design is needed to eliminate both backreflection and scattering. For optimal results, this design must be done in 3-D.

Optimisation Techniques

When we want to optimise a design for a bend or a splitter, a large parameter space is available. One can move around holes near the bend, change diameters of holes or even add arbitrary structures.

A number of the designs below are the result of the optimising software *Kallistos*, developed by Photon Design[1]. This software, which acts as a plug-in for the eigenmode expansion tool *Fimmwave*, can systematically search a parameter space of a photonic problem for the best solution [68].

This tool shows its strength in combination with 2-D simulations using eigenmode expansion (i.e. 1-D modes). However, the resulting 2-D structures should be further optimised in 3 dimensions.

Bend Designs

The simplest bend in a W1, illustrated in figure 3.13a, consists of two W1 waveguides joined together at an angle of $60°$. If we look at the field profile in the bend, we see that there is a resonance in the bend. Much of the light is reflected, and part of the light is transmitted. However, because a simple W1 supports both an even and an odd mode, not all the light is coupled into the even mode. A possible solution is to make the waveguide locally single-mode near the bend [69].

When using the bend as a cavity, transmission can be increased for the resonant wavelength. The cavity can be enhanced by moving holes around in the bend, or changing hole diameters. However, even with very small cavities the resulting bend doesn't have a very large bandwidth. Also, a resonant cavity has a very high local field, making it

[1] Photon Design web site: *http://www.photond.com*

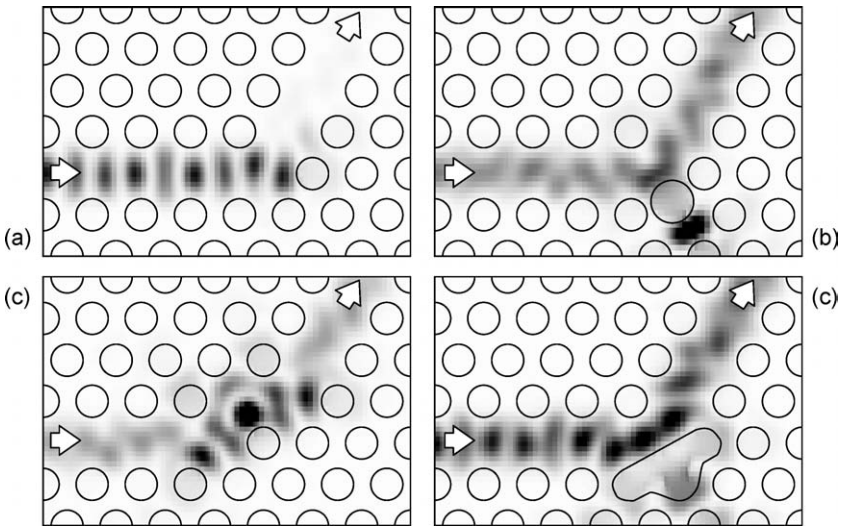

Figure 3.13: Bends in a W1 photonic crystal waveguide, with field plots at $1.55\mu m$ (a) A simple bend. (b) A bend with a resonant cavity. (c) A bend with a displaced hole. (d) A steering bend with a corner mirror. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$.

very sensitive to out-of-plane scattering. Figure 3.13b shows such a bend, with one larger, displaced hole.

It has recently been shown that simply steering the light through the bend generates a more broadband transmission [70]. Also, as the bend is no longer a resonator, the local field strength is much lower, and the bend is less likely to scatter out-of-plane. Figure 3.13c shows a more gentle photonic crystal bend where a single hole has been displaced in the bend.

While this bend has higher transmission over a broader wavelength range, it still has the problem of interfacing between two multimode waveguides. In order to couple from the ground mode into the ground mode, the bend should guide only a single mode. Therefore, just as with a straight waveguide, the optical volume of the bend should be reduced. This can be done by increasing the fraction of air in the bend. Figure 3.13d shows a bend design where a slot has been etched between two holes. This solution resembles a corner reflector in the sense that the slot also reflects the light through the bend by total internal reflection.

While these 2-D designs give very good transmission results, the final design should be studied in 3-D. In SOI, with its high vertical in-

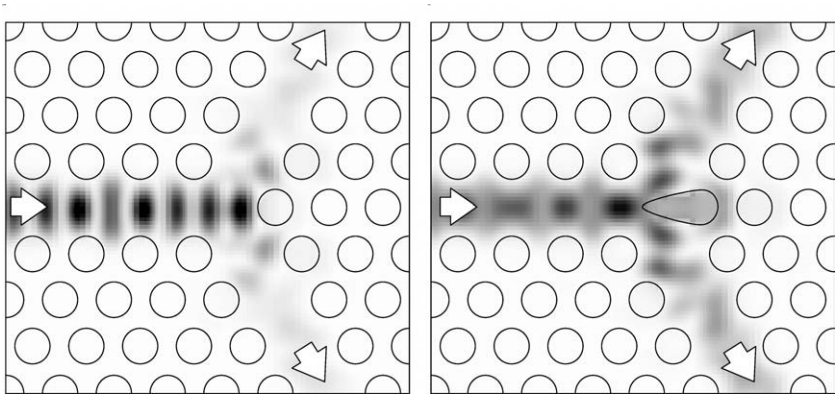

Figure 3.14: Splitters in a W1 photonic crystal waveguide, with a field plot at a wavelength of $1.55\mu m$ (a) A simple splitter. (b) A splitter with reduced volume. Lattice pitch $a = 500nm$, lattice hole diameter $\oslash = 320nm$.

dex contrast, the behaviour of a 2-D designed component could be far from optimal. This is especially true for resonant designs, where the resonance of the cavity is very much dependent on the 3-D geometry. Steering designs, like the one in figure 3.13d, are less sensitive to the vertical geometry. However, additional optimisation in 3-D is needed to minimise the out-of-plane scattering

Splitters

For splitters, a similar reasoning can be used. In figure 3.14a we see a simple splitter with three W1 waveguides joined together. This is a highly unoptimised design. For symmetry reasons alone, the splitter can only couple $44\%(4/9)$ of the light into each arm [71].

As with the bends, a suitable resonant structure can be designed to match the incoupling waveguide to the two outcoupling waveguides [72]. Of course, this solution will again be only valid within a certain wavelength range.

Alternatively, the optical volume of the splitter could be lowered [73]. Just as with the bend, a steering approach can be used to smoothly split the light and guide it into both arms [70]. Such an example is shown in figure 3.14b.

Conclusion: Bends and Splitters

Bends and splitters in photonic crystal waveguides are there to guide light. They should do this with a constant performance over a broad wavelength range. While using a resonant cavity can yield high transmission, it has a relatively narrow bandwidth. Therefore, the best designs gently steer the light. In a W1 waveguide, which supports multiple guided modes, the bend or splitter volume should be reduced to support only a single guided mode, in order to avoid exciting the unwanted guided mode. This can be done by increasing the fraction of air holes in the bend. An additional advantage is that because these designs are meant to have a large bandwidth, they are also tolerant to fabrication errors.

When we move our 2-D designs to 3 dimensions, we find that there is always a significant drop in performance due to out-of-plane scattering. Indeed, as the translational symmetry of the photonic crystal waveguide is broken, coupling to the radiative modes is possible. Further optimisation in 3-D of the design can lower these scattering losses, but it is nearly impossible to eliminate them completely.

3.2.5 Cavities: Resonators and Coupled Cavity Waveguides

When a local defect is created in a photonic crystal, e.g. by removing a single hole, a cavity is formed where light is confined in one or more bound states. Depending on the quality of the confinement, these states, or modes, exist only in a narrow frequency range.

A Waveguide and a Laterally Coupled Cavity

When we bring such a resonant defect near a photonic crystal waveguide, as illustrated in figure 3.15, light can couple back and forth between the cavity and the waveguide. However, only at the cavity resonance light will couple effectively. This way, specific wavelengths of light can be dropped from the waveguide, and the cavity acts as a wavelength filter. Simplified, the cavity resonates when the wavelength of the light in the material fits an integer number of times in the cavity. This means that the larger the cavity, the more modes are supported.

However, in this 2-D system, the only way light can escape from the cavity is by coupling to the waveguide. In such a system, with only one port, there is no net flow of light from the waveguide to the

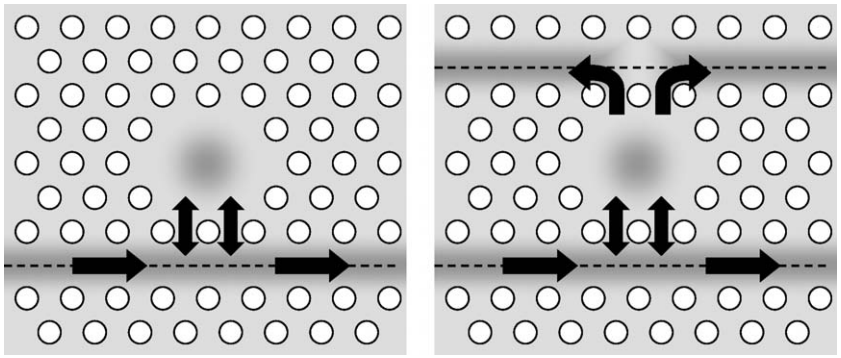

Figure 3.15: W1 Photonic Crystal Waveguides with Cavity. Left: With a single waveguide, there is no interaction as the cavity has only one port. Right: with two waveguides, the cavity can couple light from one waveguide to the other for the resonance wavelengths.

cavity. In order to work effectively as a wavelength drop filter, a second coupling mechanism is needed to couple light out of the cavity. For optimal drop functionality, the two coupling mechanisms should have the same efficiency.

An example of a cavity-based drop filter is illustrated in the right part figure 3.15. Instead of coupling the cavity with a single waveguide, it is coupled to two waveguides in an identical way. This way, the resonant wavelengths are coupled from one waveguide to the other. In this system, the coupling to the cavity is identical for both waveguides, giving rise to high coupling efficiency at the drop frequency. However, careful design is needed to tune the cavity to the correct wavelength.

The quality of a cavity is expressed by a number of parameters. The most prominent is the *quality factor* or Q of the cavity. Q is given by 2π times the number of round trips the light can make inside the cavity until its intensity is reduced by a factor of e due to the cavity losses (material losses, scattering or coupling to external waveguides) [1]. The higher Q, the sharper the transmission peak in the drop port and therefore the better the wavelength selectivity. We can derive Q from the transmission spectrum, as it is inversely proportional to the width of the transmission peaks:

$$Q = \frac{\lambda}{\Delta\lambda_{3dB}} \tag{3.1}$$

Cavities can resonate at various wavelengths, creating a comb of peaks in the drop spectrum. Typically, the larger the cavity, the closer

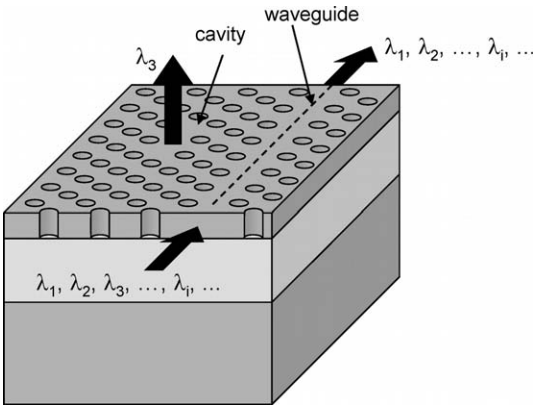

Figure 3.16: Photonic crystal cavity for use as out-of-plane drop filter. Light is coupled from the waveguide to the cavity and then vertically to a lens or fibre.

the peaks. The *free spectral range* (FSR) of the cavity is given by the wavelength difference of the resonances. When we want to drop a single wavelength channel from a broadband signal, a high FSR is needed. This way, the other resonances fall outside the operation wavelength range.

The *finesse* of the cavity combines both the wavelength selectivity as it is proportional to both Q and the free spectral range:

$$f = \frac{Q}{\lambda} FSR = \frac{FSR}{\Delta\lambda_{3dB}} \tag{3.2}$$

So, for a good wavelength-selective cavity, the cavity should be small, giving it a large FSR ($> 50nm$), and at the same time have a high Q, typically of several thousand. This last part requires that there is only weak coupling between the cavity and the waveguide.

However, when implementing this type of structure in a photonic crystal slab, light can also escape from the cavity by scattering out-of-plane. These additional losses reduce the quality of the cavity. In some applications, one can make use of this scattering loss and use the out-of-plane scattering as a drop port. This approach has been demonstrated by NTT/Kioto Univ [18], and is illustrated in figure 3.16. Here, the light at a certain wavelength is coupled out-of-plane by the cavity and can be collected by a fibre.

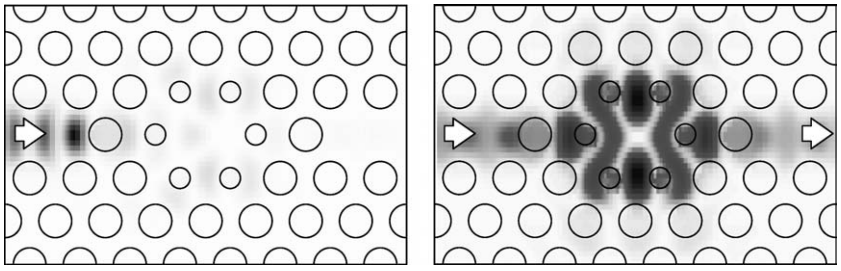

Figure 3.17: Photonic crystal in-line cavity for wavelength-selective filtering.Left: Off-resonance, right: On resonance. Only the resonant wavelength passes through the cavity, while the rest is reflected. With this geometry, Q values of up 500 to have been demonstrated [74].

In-line Cavities

Instead of using a cavity which is laterally coupled to the waveguide, one can also embed the cavity in the waveguide and use it as a wavelength selective transmission filter [16, 17]. An example, demonstrated by Glasgow University, is given in figure 3.17. Here, only the resonant wavelengths will pass through the cavity.

We can also make such a cavity in a so-called hetero-photonic crystal waveguide, i.e. a photonic crystal waveguide with step-wise change in lattice parameters in the propagation direction [75]. The change in lattice parameters can shift the waveguide in and out of cut-off, creating a mirror. Cavities in this system with a Q larger than 40000 have been demonstrated [76].

Coupled-Cavity Waveguides

When placing cavities close together, they will couple and light will be transferred from one cavity to its neighbour(s). This way, a waveguide can be created, called a *coupled-cavity waveguide* (CCW) or alternatively, a *coupled-resonator optical waveguide* (CROW) [77]. With such waveguides, sharp bends are also possible, as there is considerable freedom to position neighbouring cavities. Also, as each cavity can have one or more neighbours, depending on the crystal symmetry, splitters are as easy to implement as an ordinary waveguide[78]. This is illustrated in figure 3.18.

However, when implemented in a photonic crystal slab, CCWs have a considerable drawback. A straight CCW has a periodicity with a period which is several times the lattice constant a of the photonic crystal.

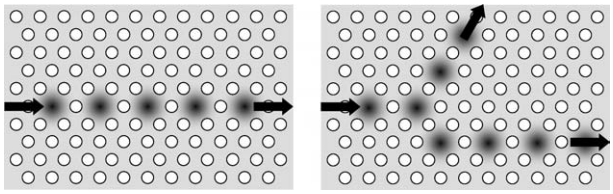

Figure 3.18: Coupled-cavity waveguide. (a) straight waveguide composed of H1 cavity (one hole removed), (b) splitter and bends.

Hence, the first Brillouin zone is several times smaller than that of the photonic crystal. When operating in the PBG of the photonic crystal with periodicity a, the CCW's guided modes have already been folded several times and are well above the light line. Therefore, CCW modes will always couple light to the radiation continuum.

In a 3-D photonic crystal, CCWs do offer a practical waveguiding solution, as light is confined by the 3-D photonic crystal. This has already been demonstrated extensively for microwave and far-IR wavelengths [79, 80].

3.2.6 Dispersion and Slow Waves

Photonic crystals have very strong dispersion. Even though this makes simple waveguiding difficult, it can enhance other functions. We will briefly discuss a few of them, without going into too much detail.

Angular Dispersion and the Superprism

Instead of using a defect to create a functional element, we can also exploit the special properties of the lattice. We have seen in chapter 2 that the dispersion of the photonic crystal differs for the various propagation directions in the lattice. This results in a number of interesting applications.

In a superprism we see that an incident beam with multiple wavelengths is split up in the photonic crystal into its constituent wavelengths [19]. This effect is up to 1000 times stronger than in an ordinary wavelength-splitting glass prism. This is caused by the irregular shape of the dispersion surface near the band edge. When multicolored light hits the interface at an angle, the various wavelengths feel a material with a very different dispersion relation, and are refracted under dif-

ferent angles. The effect has also been demonstrated in a 2-D photonic crystal slab [81, 82].

Active Functions

We have seen that it is not straightforward to achieve broad bandwidth in a photonic crystal waveguide. In a simple W1 waveguide, for instance, the bands below the light line are very flat. While this is not well suited for pure waveguiding, flat bands can be useful, because of their low group velocity.

In regions with low group velocity, light pulses travel slowly in the photonic crystal. This can be extremely useful for active functions, because the interaction time with the material is much longer for a given length. Also, because photonic crystal waveguides have very small cross sections, the intensity of the light is very high, making it suitable for nonlinear applications.

For example, when we have a photonic crystal lattice with gain, light near the band edge travels very slowly, and is therefore continuously amplified by bouncing back and forth through the crystal. At the band edge, where the group velocity is zero, this effect is strongest, and lasing can occur. This principle, on which the operation of a *distributed feedback laser* is based [20] was recently demonstrated in 2-D photonic crystals with high contrast [21, 22].

For nonlinear functions, photonic crystals can be used to make compact optical switches [83] and flip-flops using the Kerr effect [25, 26]. The periodic structure can also be used to create a waveguide in which a wave has the same phase velocity as a wave with double frequency. Using this phase matching, a very efficient way of generating light at the second harmonic is possible when the material has a sufficient $\chi^{(2)}$ value [23]. This, in turn, can be used for wavelength conversion [24].

The use of photonic crystal structures for active functions is becoming an increasingly important part of photonic crystal research. However, Silicon-on-insulator is not ideally suited for this purpose. Therefore, in this work, we will further focus solely on passive waveguiding devices.

3.2.7 Conclusion

We have discussed a number of waveguide designs and components in photonic crystals. For waveguiding purposes, we see that each design

has its strengths, but we have found no design that satisfied all our criteria. In the SOI system, broad bandwidth is especially hard to achieve because of the low light line of the oxide cladding. Also, for some theoretically good waveguide designs the robustness with respect to fabrication accuracy and sidewall roughness is questionable.

For functional components, photonic crystals carry more promise. When properly optimised in 3-D, sharp bends and abrupt splitters can by reasonably efficient. Also, small cavities with relatively high quality are possible.

However, the unique strength of photonic crystals lies in its very strong dispersion properties. Especially slow wave effects could prove useful for active functionality.

3.3 Photonic Wires

The principle of photonic wires is the same as of conventional optical waveguides: light is confined in a narrow core of high index material surrounded with a cladding of lower index material. For photonic wires, the index contrast between core and cladding is very high. This gives rise to very strong confinement, which makes it possible to make very sharp bends without radiation losses in the bend. However, there is no photonic bandgap to stop light from radiating away once it escapes the confinement of the wire.

Photonic wires are not periodic structures, and therefore their dispersion relation is far less exotic than photonic crystals. This makes them more predictable and easy to design. Moreover, photonic wires are broadband with a fairly linear dispersion, making them very suitable for simple waveguides.

3.3.1 Straight Waveguides

The principle of photonic wires is based on strong lateral confinement. In the vertical direction, a high refractive index contrast is not strictly necessary. In this work, we will focus on the application of photonic wires in Silicon-on-insulator, where there is very strong vertical confinement. Just as with photonic crystal structures, we have the option to etch deeply or just etch the core layer, as illustrated in figure 3.19. For photonic wires, deep etching can slightly reduce the bend losses, as the bottom cladding now has a lower refractive index.

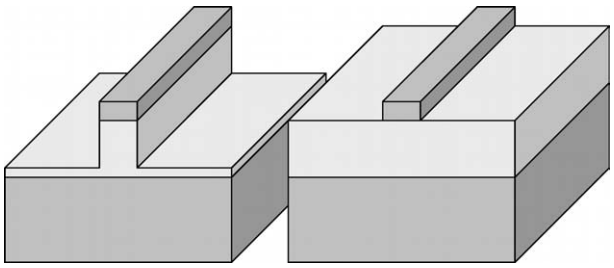

Figure 3.19: SOI photonic wires with deep and shallow etch.

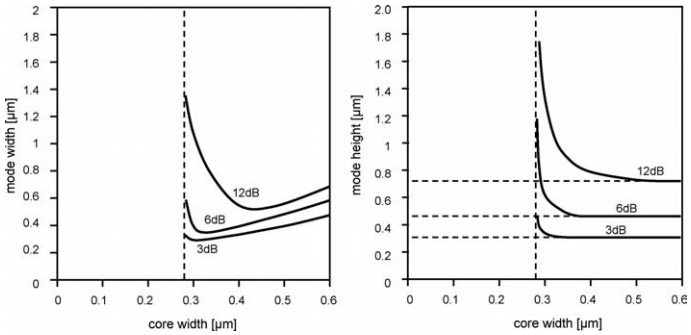

Figure 3.20: Mode height and width of the guided mode of an SOI photonic wire as a function of core width. The $3dB$, $6dB$ and $12dB$ mode width is plotted. Left: Mode width. Right: mode height.

In order for a photonic wire to function correctly, it needs to have only a single guided mode. Because photonic wires use a high index contrast, a narrow waveguide is needed to fulfill this condition. In SOI with a core thickness of $220nm$, a photonic wire becomes single mode for the TE polarisation at a core width of about $550nm$ for wavelengths around $1.55\mu m$. If the waveguide gets broader, more guided modes are supported. When the waveguide symmetry is broken, such as in bends, the waveguide modes can mix, which is undesirable.

The core dimensions also influence the shape of the guided ground mode. For a wide core (i.e. a multimode waveguide), the mode is very strongly confined in the vertical direction, closely resembling the slab mode. When the waveguide gets narrower, so does the mode. However, when the waveguide gets too narrow, most of the mode energy is outside the waveguide core and the mode becomes broader again. This is illustrated in figure 3.20.

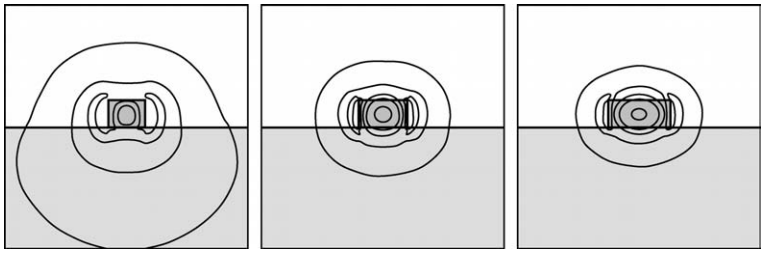

Figure 3.21: Mode profiles for the ground TE mode of a photonic wire a core width $w_{core} = 0.3, 0.4, 0.5\mu m$. The core thickness $d_{core} = 0.22\mu m$, and the oxide cladding is supposed to be infinitely thick. When the wire gets narrower, the mode extends more in the vertical direction.

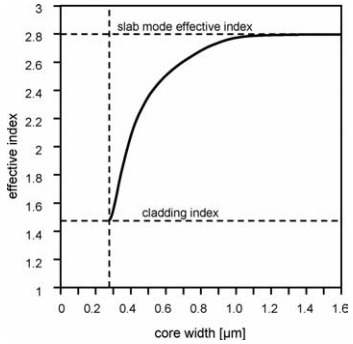

Figure 3.22: Effective index of the fundamental photonic wire mode as a function of core width w_{core}. Because of the asymmetric vertical cladding the waveguide has a cut-off width, below which no guided mode is supported.

We can see a number of mode profiles for narrow wires in figure 3.20. The core width also has an effect on the vertical confinement. When the core becomes narrower, the mode will extend further into the top and bottom cladding. As will be further discussed in chapter 4, this has a serious effect on the substrate leakage.

Because the waveguide cladding is asymmetric (oxide below and air on top), its dimensions cannot be reduced indefinitely. When w_{core} gets too small, all the light is forced into the bottom cladding and the waveguide does not support a guided mode any longer. For a core thickness $d_{core} = 220nm$ the cut-off width is $w_{core} = 285nm$. The effect of the core width on the effective index is plotted in figure 3.22.

Because photonic wires confine the light only by refractive index contrast, they are very sensitive to scattering, both in-plane and out-of-plane. Unlike in photonic crystal slabs, there is no photonic bandgap to

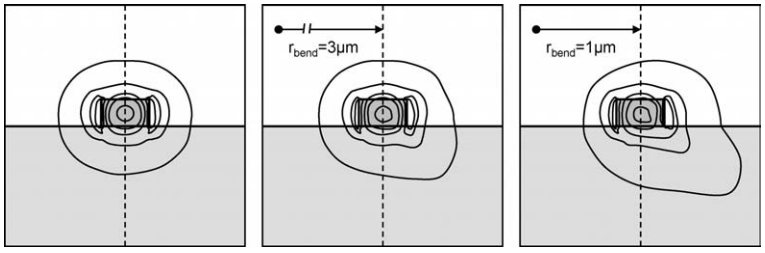

Figure 3.23: Bends modes of photonic wire bends with different bend radii r_{bend}. Left to right: $r_{bend} = \infty$ (straight waveguide), $r_{bend} = 3\mu m$,$r_{bend} = 1\mu m$.

stop the light from radiating away into the cladding. Therefore, photonic wire propagation losses are a good measure for the fabrication quality of nanophotonic structures in general.

3.3.2 Bends and Mirrors

Because of the strong confinement, it is possible to make sharp bends in a photonic wire without radiating too much light. In a bend, the waveguide mode will change with respect to the mode of the straight waveguide. This is because the wave fronts at the outside of the bend have to propagate faster than at the inside of the bend. This effect moves the power of the mode towards the outside of the bend [1]. This effect gets stronger when the bend radius is smaller. This is illustrated in figure 3.23.

In conventional waveguides with low confinement, this causes a mode mismatch at the interface between the straight waveguide and the bend. Therefore, the axis of the bent waveguide is always slightly offset in order to minimise this mismatch [84]. However, because of the strong confinement in photonic wires, the center of the mode is hardly shifted towards the outside of the bend. Therefore, such an offset is not needed in photonic wires.

The faster travelling wave fronts at the outside of the bend are limited to the speed of light. Therefore, at a certain distance outside the waveguide core, the light will no longer follow the bend, but radiate away. Again, this effect gets stronger as the bend radius gets smaller. In conventional waveguides, this limits the reduction of the bend radii to $100\mu m$ or more. In photonic wires, however, the strong confinement allows for much smaller bend radii. Figure 3.24 shows a 2-D simulation of a photonic wire bend with a bend radius $r_{bend} = 3\mu m$.

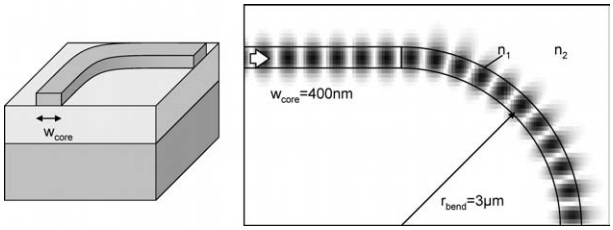

Figure 3.24: A 2-D simulation of a bend in a photonic wire. $n_1 = 2$, $n_2 = 1$, $w_{core} = 0.4\mu m$, $r_{bend} = 3\mu m$.

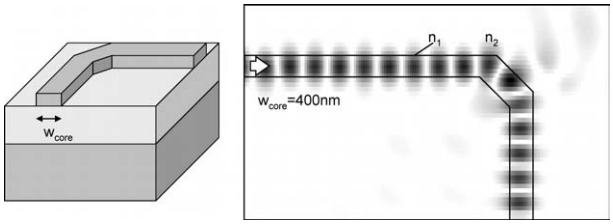

Figure 3.25: A 2-D simulation of a corner mirror in a photonic wire. $n_1 = 2$, $n_2 = 1$, $w_{core} = 0.4\mu m$. It has a transmission of 75%.

Instead of using a gradual bend, one can use the high refractive index contrast to make a mirror that deflects the light from the waveguide to another waveguide in another direction [27, 28, 29]. The principle is illustrated in figure 3.25. We can see that a large amount of light is deflected by the mirror into the other waveguide. However, while this component is more compact than a gradual bend, it is more prone to losses. The abrupt discontinuity can couple light from the wire to radiation modes, just as with a photonic crystal bend. Also, it is not straightforward to fabricate a perfectly smooth mirror.

3.3.3 Splitters, Couplers and Intersections

As photonic wires are very similar to conventional waveguides, it is easier to port concepts for devices than for photonic crystals. However, many designs are adiabatic, i.e. the structure changes slowly along the propagation direction so the waveguide mode can adapt itself continuously. This means that the device is quite long, which, on first sight, makes it less suitable for nanophotonic applications. Still, due to the strong confinement, these adiabatic designs can in some cases be shrunk considerably.

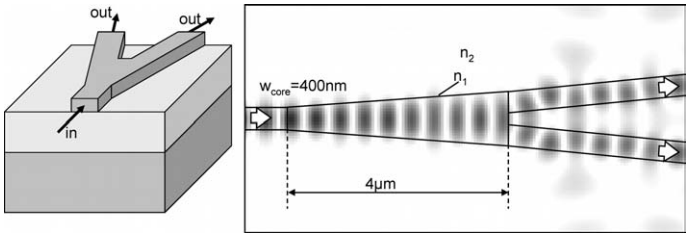

Figure 3.26: A 2-D simulation of a simple splitter in a photonic wire. $n_1 = 2$, $n_2 = 1$, $w_{core} = 0.4\mu m$. The splitter length is only $4\mu m$, with 95% transmission.

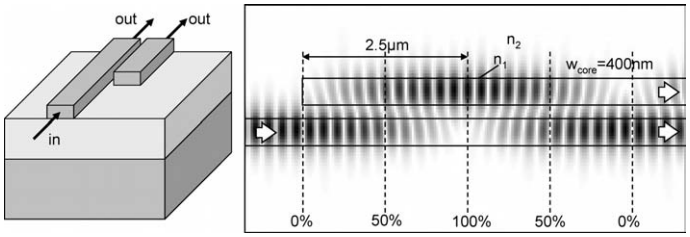

Figure 3.27: A 2-D simulation of a directional coupler in a photonic wire. $n_1 = 2$, $n_2 = 1$, $w_{core} = 0.4\mu m$, $w_{trench} = 0.2\mu m$. All power is coupled from the first waveguide to the next over $2.5\mu m$.

Figure 3.26 shows a simple $3dB$ splitter, where the waveguide is widened gradually and then split in two. The splitter is only $4\mu m$ long, but still has a transmission of over 95% over a wide wavelength range. With further optimisation of the splitter geometry, the transmission can be increased.

Directional couplers are also possible, by positioning 2 photonic wires together, as in figure 3.27. When the wires are sufficiently close together, the coupling length required to couple all the power from one waveguide to the next can be very short. For the component in figure 3.27, the coupling length is only $2.5\mu m$ with a waveguide spacing of $200nm$. However, this strong coupling makes the component very sensitive to fabrication tolerances, and also wavelength dependent. A small change in the waveguide width or spacing will drastically influence the coupling.

With the right length, a directional coupler can be used as a splitter with arbitrary splitting ratio. Figure 3.28 shows a $3dB$ splitter based on the coupler from figure 3.27. Note that there is also some coupling in the splitter section where the waveguides are still close together.

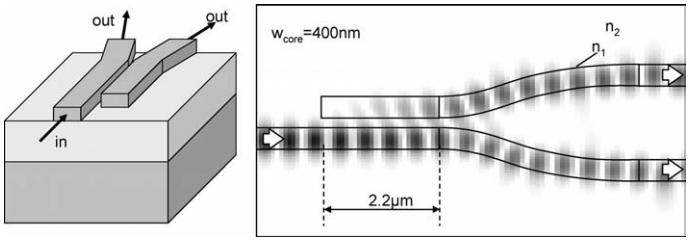

Figure 3.28: A 2-D simulation of a $3dB$ splitter with a directional coupler in a photonic wire. $n_1 = 2$, $n_2 = 1$, $w_{core} = 0.4\mu m$, $w_{trench} = 0.2\mu m$.

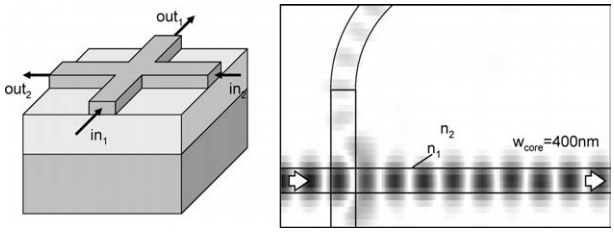

Figure 3.29: A 2-D simulation of a crossing of two photonic wires. $n_1 = 2$, $n_2 = 1$, $w_{core} = 0.4\mu m$. The crosstalk is only 2%, and the transmission is over 93%.

The large index contrast and the narrow width make that the guided mode has components at large angles, both vertically and in-plane. In other words, the waveguide has a large numerical aperture. In spite of that, it possible to cross two waveguides at right angles with very little crosstalk. Of course it helps that the size of the crossing is equally small, due to the small waveguide width. In figure 3.29 there is only 2% crosstalk and 1% reflection at the intersection. Over 93% of the light continues its way in the original waveguide.

3.3.4 Ring and Racetrack Resonators

Because short bend radii are possible in photonic wires, one can make very compact ring resonators [30]. Figure 3.30 illustrates a ring resonator and a racetrack resonator. In the latter case, the coupling section between the straight waveguide and the ring is longer, so higher coupling efficiency is possible.

A ring or racetrack resonator is a travelling-wave resonator [85]. This means that there is no reflection in the cavity, but the ring resonates when the wavelength of the light fits an integer number of times in the

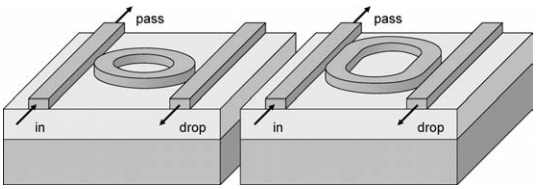

Figure 3.30: A ring resonator and a racetrack resonator. In the racetrack resonator, the coupling is enhanced by the additional straight coupling section.

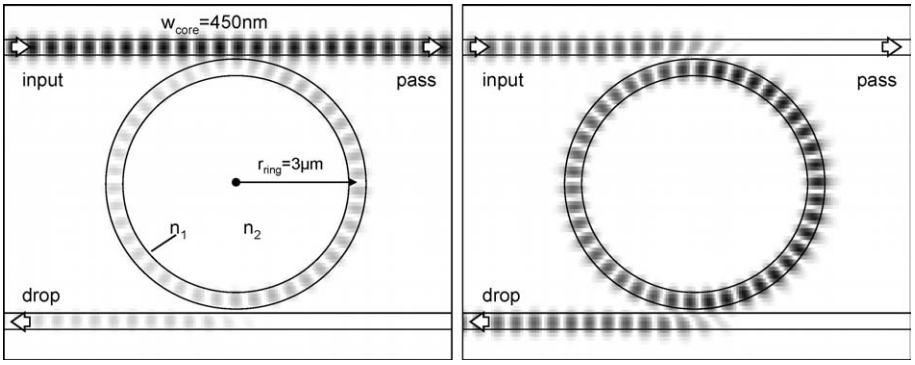

Figure 3.31: A 2-D simulation of a ring resonator in a photonic wire. Left: Off-resonance ($\lambda = 1555nm$). Right: On resonance ($\lambda = 1530nm$). $w_{core} = 400nm$, $w_{gap} = 150nm$, $r_{ring} = 3\mu m$.

circumference of the ring. When the ring resonates, light is coupled from the input waveguide to the output waveguide. Figure 3.31 shows the electromagnetic fields in the ring on and off the resonance wavelength. The transmission in both waveguides as a function of wavelength is plotted in figure 3.32.

Ring resonators can reach high Q values. However, for optimal coupling efficiency at the resonance wavelength, the ring losses should be kept to a minimum, and the coupling between the ring and both waveguides should not differ too much. Only at a critical coupling ratio that compensates exactly for all cavity loss, all power is extracted from the input waveguide at resonance and crosstalk is reduced to zero [86]. Because the coupling is strongly dependent on the width of the trench between the ring and the straight waveguides, the fabrication tolerances of these structures are very strict, typically of the same order as the hole diameters of photonic crystal holes.

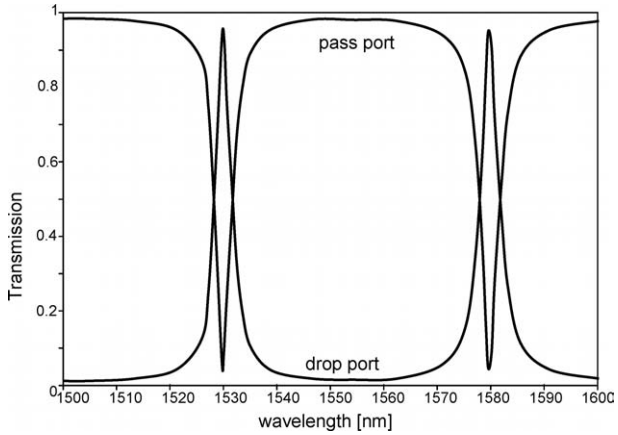

Figure 3.32: Transmission spectrum of the ring resonator in figure 3.31.

3.3.5 Conclusion

Photonic wires are very good for waveguide components. They have a broader bandwidth and less wavelength sensitivity than their photonic crystal counterparts. Also, it is possible to make sharp bends. Wavelength-selective components in wires can be based on ring resonators. However, these are just as sensitive to fabrication accuracy as photonic crystal cavities.

3.4 Fibre Couplers and Mode Converters

While photonic crystals and photonic wires make it possible to reduce the size of on-chip waveguide components, the reduced size makes the coupling of light to and from optical fibres far from easy. Figure 3.33 shows the waveguide mode of a $450nm$ photonic wire on the same scale as the core of a single-mode optical fibre. When both waveguides are butt-coupled, the coupling efficiency is of the order of 0.1%.

The most obvious method for increasing the coupling efficiency is by using a lens or lensed fibre to focus the light onto the nanophotonic waveguide. However, this approach is very sensitive to alignment errors, because one tries to align two areas with submicron dimensions with respect to each other. The obvious way of reducing the alignment tolerances is by adding a spot-size converter to the chip, which dramatically increases the spots that need to be aligned. While it is relatively easy to make an adiabatic taper that broadens the mode, it is

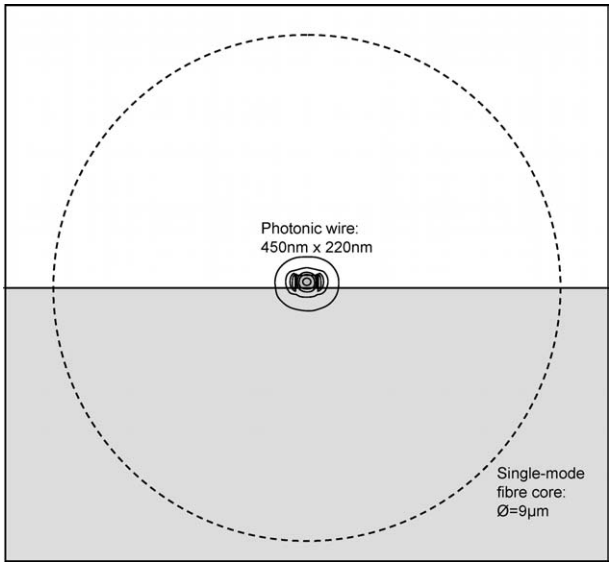

Figure 3.33: Comparison between a single-mode fibre core and a nanophotonic waveguide mode on the same scale.

not straightforward to change the mode's vertical dimension. Some solutions include the application of an additional polymer optical waveguide on top of the semiconductor waveguide [31], or alternatively, the use of a multi-layer semiconductor stack to make a broad ridge waveguide underneath the nanophotonic waveguide [32, 33].

While both solutions offer satisfactory coupling efficiencies, they use an adiabatic approach. As this requires a long propagation length, these tapers consume a large chip area, cancelling some of the gains of nanophotonic waveguides.

Therefore, an alternative approach is needed. In this section, we will briefly discuss the use of a strong periodic grating to couple the light from an optical fibre into an broad SOI waveguide. The broad waveguide is then coupled to a narrow photonic wire using a short, irregular structure to image the guided mode of the broad waveguide to the narrow waveguide.

3.4.1 Grating-Based Fibre Couplers

Instead of coupling light from a fibre to a thin SOI waveguide using a taper, we can use a 1-D periodic grating to couple light to a fibre per-

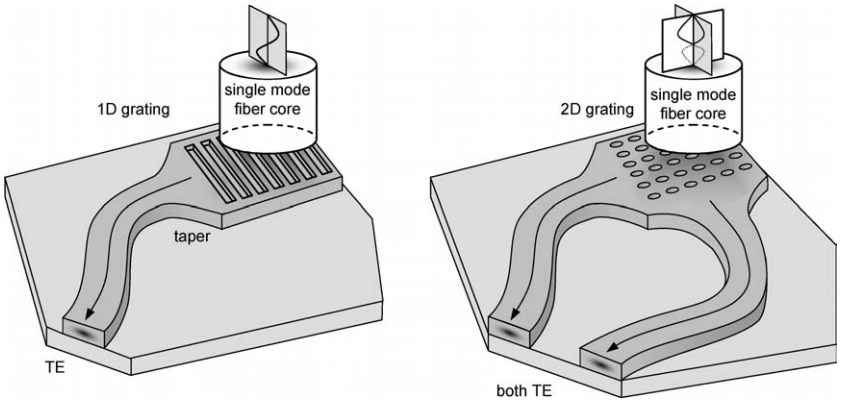

Figure 3.34: Principle of a photonic-crystal based fibre coupler. Light is coupled by a strong grating from the vertically oriented fibre to a thin, but high-contrast waveguide. Left: A 1-D grating couples only one fibre polarisation. Right: a 2-D grating couples both fibre polarisations into their own waveguide.

pendicular to the surface[2]. The principle is illustrated in the left part of figure 3.34. The use of gratings for the outcoupling of light is not new. In most cases, however, a very weak grating is used, which requires a length between $100\mu m$ and $1mm$ [87] to create a beam focused to a detector or fibre several mm above the structure. For nanophotonic applications, the coupling structure should be much smaller. Therefore, the fibre is positioned close to the grating, and the grating is made stronger by etching deeper into the Silicon [34].

Because we want to couple vertically, the grating is a second-order grating of which we use the first order of diffraction. Therefore, we need a larger period compared to the photonic crystal lattices discussed earlier in this chapter. Typically, we used periods of $580 - 600nm$. In contrast with planar photonic crystals, the grating is not etched through the core, but only $50nm$ deep. Figure 3.35 shows a simulation of such a structure. 35% of the power is coupled upward, 40% is coupled downward. The coupling efficiency to the fibre mode is 28%. For the simulated wavelength, the light is coupled out at an angle slightly off the normal direction.

This fibre coupler needs a strong interaction with the grating. Also, a strong vertical index contrast is required for high efficiency. This

[2]The research on the grating-based fibre coupler has been carried out by Dirk Taillaert in the context of his PhD.

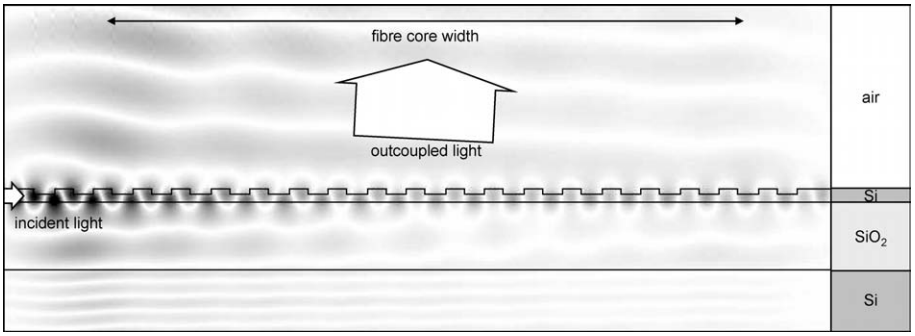

Figure 3.35: Simulation of a 1-D fibre coupler. The theoretical coupling efficiency to fibre is 28%.

makes SOI ideally suited for this purpose. However, the grating couples light from the planar waveguide both upward and downward. Therefore, the oxide thickness should be chosen carefully. This way, the additional reflection of the downward wave at the interface between the oxide buffer and the Silicon substrate can be made to interfere constructively with the upward wave.

The use of such a 1-D periodic grating has a considerable drawback: only one polarisation from the fibre is coupled to the TE polarisation of the planar waveguide. The other polarisation is not coupled at all. Because the light in most standard telecom fibres can have any polarisation, the coupling efficiency is completely unpredictable in this case.

However, by creating a 2-D periodic grating, both fibre polarisations feel the periodicity, and will be coupled to a planar waveguide. This concept is illustrated in the right part of figure 3.34. Because both fibre polarisations are coupled into an identical waveguide, and both in the TE polarisation. This way, both fibre polarisations can be processed in an identical circuit, and at the end, joined together by the same type of structure [35].

3.4.2 Compact Spot-size Converters

The grating-based fibre couplers couple light from a fibre into a $10\mu m$ broad waveguide. Therefore, we still need a taper or spot-size converter to focus the light into a nanophotonic waveguide. While this can be done by simple adiabatic tapers, these structures are still very long compared to other nanophotonic components.

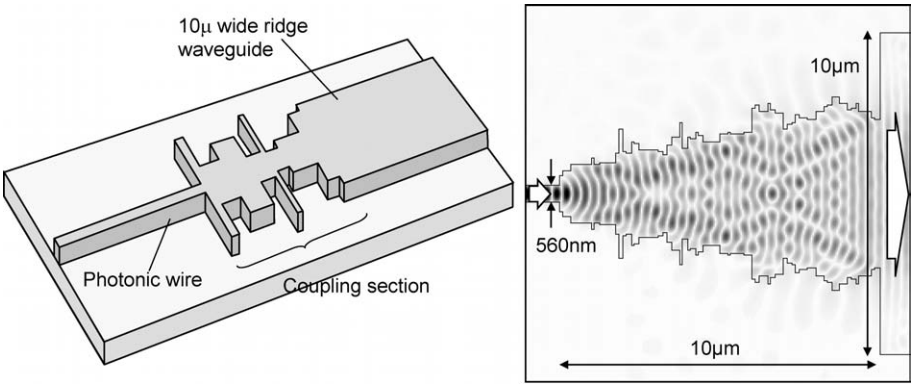

Figure 3.36: Principle of in interference-based spot-size converter. In this highly-multimode structure the interferences between the different modes can couple the light from a broad waveguide into a narrow waveguide over a short distance. Left: A typical structure. Right: a simulation result, yielding 92% transmission.

Figure 3.36 gives an alternative solution[3]. Instead of using a slowly varying taper that gently converts the waveguide ground mode without exciting the other waveguide modes, this approach deliberately excites a large number of waveguide modes which interfere heavily over a short distance [88]. By choosing the parameters correctly, this interference can give rise to a focused spot which matches the mode of the nanophotonic waveguide. As the light is also coupled from the ground mode of the broad waveguide to the ground mode of the photonic wire, the component works in both directions, due to reciprocity.

The example in figure 3.36 couples light from a $10\mu m$ wide waveguide to a $500nm$ wide photonic wire over a distance of $10\mu m$. The coupling efficiency, mode to mode, is 92%.

3.5 Conclusion

Nanophotonics offer a significant reduction in size compared to conventional photonic waveguide components. However, the reduced size and the high index contrast create a number of design issues. While photonic crystals offer a wide selection of strong dispersive properties, it is not straightforward to design a simple waveguide. On the other

[3]The research on the integrated interference-based spot-size converter has been carried out by Bert Luyssaert in the context of his PhD.

hand, the dispersive properties, together with the strong confinement, makes the structures very promising for wavelength-selective functionality.

For simple waveguides, a photonic wire is much better suited. Its behaviour is very predictable, and it is simple to make compact elementary waveguide components, like bends and splitters. However, as we will see in chapter 4, wires are sensitive to sidewall roughness, more so than photonic crystals. Therefore, very good fabrication quality is needed. This is also true for ring resonators in photonic wires. An accurate control of the ring parameters is needed to make efficient rings.

To interface nanophotonic circuits to the outside world, fibre couplers and spot-size converters with a small footprint are needed. The grating-based fibre couplers, which can also include polarisation splitting functionality, are good candidates, especially when combined with a short spot-size converter, like the interference-based designs shown here.

Overall, it is clear that for nanophotonic structures the design rules for conventional photonics have to be rethought or abandoned. Full 3-D simulations are needed to optimise a structure, and high-end technology to reliably fabricate the optimal design.

Chapter 4

Losses in Nanophotonic Waveguides

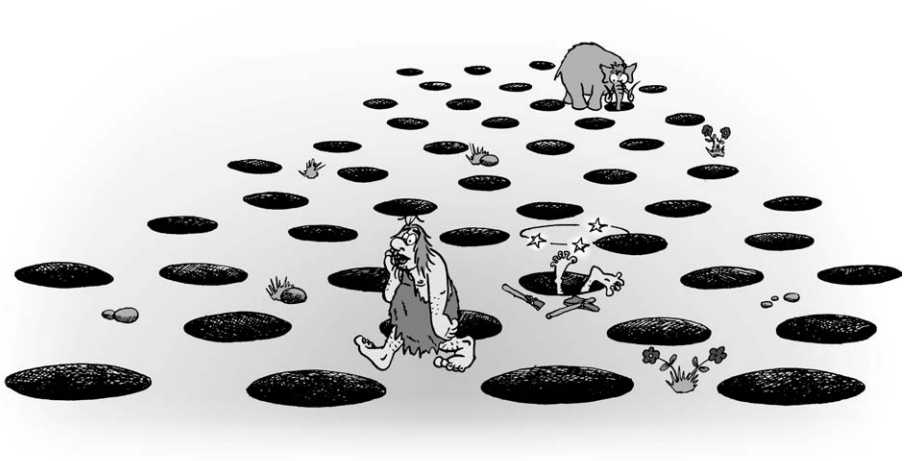

"Ug! We're lost. Where are you?"

Chapter 4

Losses in Nanophotonic Waveguides

Nanophotonic waveguides, with their high index contrast, allow very compact waveguide designs. However, this is only an advantage when the losses are sufficiently low to allow integration of multiple components onto a single chip.

Photonic wires keep light confined by a high index contrast, and this allows for very sharp bends. Photonic crystal waveguides get their unique properties from their periodicity. Through the photonic bandgap effect, they promise perfect confinement of light. However, in the real world, nanophotonic waveguides, and in particular 2-D photonic crystal slab waveguides have a number of possible loss mechanisms. In this chapter we will determine the favourable regimes for efficient waveguiding, and look where the losses come from.

4.1 Loss Mechanisms

As explained in the previous chapter, photonic crystal slab waveguides guide the light in-plane by the photonic bandgap effect, while out-of-plane, a conventional slab waveguide concentrates the light into the layer with a high refractive index. While in first approximation these two guiding mechanisms can be treated independently, the photonic crystal slab waveguide should be treated as a 3-D system.

When taking a simplistic view of the situation, the holes in the photonic crystal remove the vertical index contrast of the slab waveguide. Because of this, the light in the holes is no longer confined in the ver-

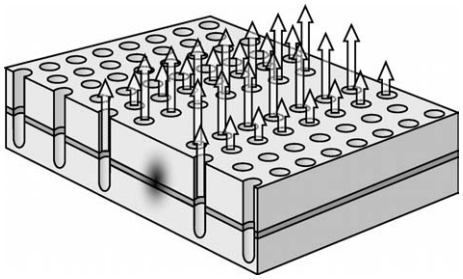

Figure 4.1: Out of plane scattering losses. Light is scattered out of plane because of the lack of vertical confinement in the etched holes.

tical direction and can scatter outward, as illustrated in figure 4.1. As the light propagates through the waveguide, this scattering introduces losses. Therefore, even a perfect photonic crystal slab is not necessarily a lossless waveguide, although we will see that this is indeed possible, and that one can reduce or even eliminate the intrinsic out-of-plane scattering losses.

In real-world structures fabricated by etching holes into semiconductor (see chapter 5), it is unavoidable that a certain amount of roughness is introduced at the sidewalls of the etched structures. Light propagating through the waveguide can scatter at such irregularities and introduce additional waveguide losses.

One important issue in the world of photonic crystal slabs is the choice of layer structure one wants to use for the vertical confinement. In effect, there is a choice between high or low vertical index contrast. When these studies were undertaken, it was by no means clear which material system performed best with respect to out-of-plane scattering losses. For high-index contrast, the argument that the confinement was better was countered by the fact that the guided mode contained components with large out-of-plane angles. On the other hand, for low index contrast, the guided mode is very much stretched out like a plane wave and therefore less confined, but has little components out-of-plane.

We have compared both types of layer structures with respect to intrinsic out-of-plane scattering and scattering at roughness, and have identified the strengths and weaknesses of both, and the regimes where low-loss operation is possible.

For these loss calculations, we assumed a slab waveguide cladding of infinite thickness. In real life, this is not the case, but the bottom cladding is a low-index layer of finite thickness on top of high-index

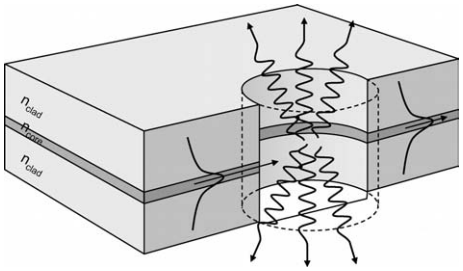

Figure 4.2: In-plane light crossing an etched hole is scattered. This can be modelled by a dipole source in the center of the hole.

substrate. In SOI, the substrate material is identical to the top waveguide core. If the bottom cladding is too thin, light will leak out of the core to the substrate. For nanophotonic waveguides, this is even more the case, because the light is stretched out vertically as the waveguide core gets narrower in the horizontal direction. At the end of this chapter, we will determine the conditions needed to reduce these leakage losses.

4.2 Out-of-Plane Scattering Losses

In photonic crystal slab waveguides, light is guided vertically by the core layer with the higher index [6]. However, in the holes where the core layer is etched away, this vertical confinement is broken, and light can scatter out-of-plane [36]. While this is a simplified view of the loss mechanisms, out-of-plane scattering is a major cause of loss in many photonic crystal slab structures.

As discussed in the previous chapter, it is possible to choose between a number of material systems with either a high or a low refractive index contrast. For many years, it has been unclear whether either low or high contrast was the most beneficial with respect to out-of-plane scattering.

A first attempt at modelling the out-of-plane scattering in the holes of photonic crystals was undertaken by Benisty [36]. He modelled each hole individually as a scattering dipole excited by the electromagnetic field in the hole. The principle is illustrated in figure 4.2.

To calculate the amount of scattered light, Benisty uses a perturbation method, starting from an analytically solvable structure which has a mode that was perfectly confined. This structure, however, did still

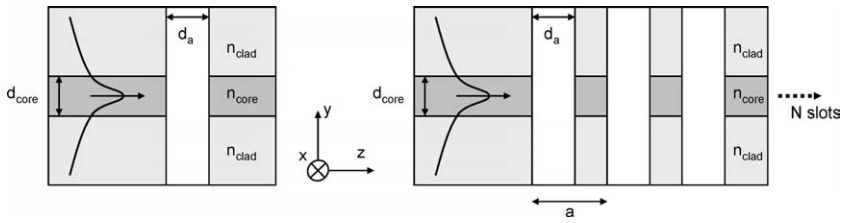

Figure 4.3: 3-layer slab waveguide with air slots. Left: Single air slot. Right: multiple periodic air slots.

have a vertical index contrast in the holes. Therefore, the index change in the holes was added to the model as a perturbation, and the subsequent scattering of light by this perturbation could be approximated as a dipole.

Benisty found that the scattering is proportional to the square of the vertical contrast in relative dielectric constant $\Delta\epsilon = n_{core}^2 - n_{clad}^2$ [37]. However, because he used a perturbation model, this approach was only valid for low refractive index contrasts [38]. Also, he assumed that the holes scattered incoherently, while in a photonic crystal the periodic stacking of holes can give rise to interference that either enhances or reduces the scattering.

4.2.1 Simulated Structure

In order to get a more complete view of out-of-plane scattering in photonic crystal slabs, we modelled the scattering numerically. For this purpose, we used the CAMFR eigenmode expansion tool. However, because the simulation of the full 3-D structure of a 2-D photonic crystal slab was not possible at the time, we modelled the out-of-plane scattering in a 2-D equivalent: a 1-D periodic arrangement of air slots in a 3-layer slab waveguide (figure 4.3). Just like a photonic crystal waveguide, this structure has a 1-D periodicity. However, in this structure, light has to cross the air slot each time, while in a photonic crystal waveguide a large fraction of the light will flow around the holes. The 1-D array of air slots will therefore have much more losses than a real-life photonic crystal waveguide, but it can give us a qualitative view of the scattering mechanism.

We use this structure as an approximation of a slab perforated with holes. This type of photonic crystal has a PBG for the TE polarisation, i.e. with the electric field in the plane of the structure, so we will use the

same polarisation for our structure. This means that the electric field E is oriented perpendicular to the simulated structure in figure 4.3.

We simulated both the situation with a single air slot in the slab waveguide, in order to compare our results against Benisty's analytical model, and the situation with a periodic stacking of slots, to study the effect of the photonic bandgap on the out-of-plane scattering. As we will see, the PBG has a profound effect.

As semiconductor appears to be the most popular candidate for 2D photonic crystals, we tried to approach common semiconductor layer structures. For the sake of simplicity, we chose a three-layer slab waveguide (top cladding, waveguide core and bottom cladding) with a core thickness d_{core} of $225nm$ and a refractive index n_{core} of 3.5. The slab waveguide of figure 4.3 is invariant in the x-direction perpendicular to the page. We chose the slab waveguide to be single mode at a wavelength of $1.55\mu m$ for all values of the cladding index n_{clad}, and the guided mode of the slab waveguide propagates in the z-direction. We then varied the waveguide cladding index n_{clad} to study the behavior with different refractive index contrasts. This change in index contrasts obviously alters the dispersion relation of this waveguide and also the band structure of a periodic structure etched into the waveguide. In order to compare our results to those of Benisty, we will write the refractive index contrast in the vertical direction as $\Delta\epsilon = n_{core}^2 - n_{clad}^2$. Some values correspond with real layer structures. The upper limit of 11.25 corresponds to a silicon membrane with air cladding. A silicon core in an oxide cladding, not unlike SOI, has a value of $\Delta\epsilon = 10$, while a low-contrast $GaAs$-$Al_xGa_{1-x}As$ heterostructure has a $\Delta\epsilon$ of about 2.

To measure the out-of-plane scattering, we launch light from the left into this single-mode waveguide and measure the reflection and transmission into the waveguide mode. As only the light in the single waveguide mode is guided, the remainder of the light is lost due to radiation. Eigenmode expansion is ideally suited for this type of simulations, as it can easily extract the power associated with a specific waveguide mode, be it leaky or guided. For these simulations, we used CAMFR because perfectly matched layers are well supported, and in order to correctly simulate radiation losses, these absorbing boundary conditions are necessary.

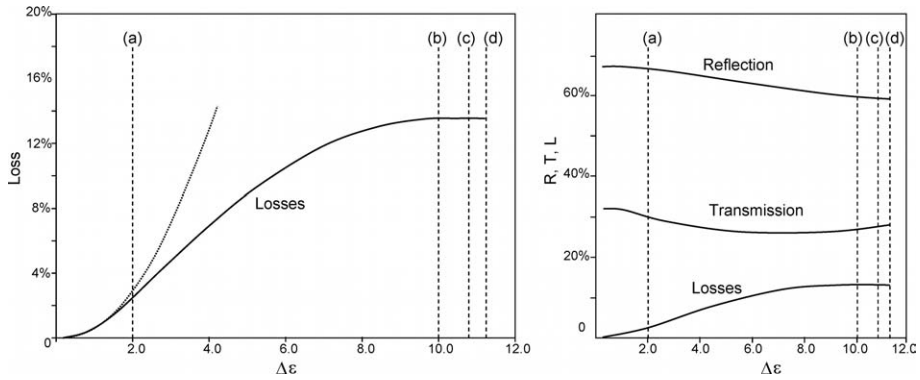

Figure 4.4: Out-of-plane scattering losses for a single air slot with a width $d_a = 0.28\mu m$. For low $\Delta\epsilon$ the losses increase with the square of $\Delta\epsilon$ (indicated by the dotted curve), but they level out at high index contrasts. The right part of the figure shows the reflection and transmission of the guided mode.

4.2.2 Scattering at a Single Air Slot

When we limit the problem to the scattering of a single air slot in the slab waveguide, we can treat this as the incoherent limit where each hole radiates independently of the other holes, and scattering losses of individual holes or slots are added up. In this approximation we disregard the effect of the periodic structure on the scattering. Figure 4.4 shows the losses of a single air slot with a width d_a of $0.28\mu m$ at a wavelength $\lambda = 1.55\mu m$. Reflection and transmission of the guided mode are also plotted. Simulations of the losses of a single air slot confirm the notion that out-of-plane scattering increases with higher refractive index contrast. For very low index contrast the losses increase with the square of $\Delta\epsilon$, as predicted by the model of Benisty [37]. Moreover, the radiation of a single air slot is very similar to a two dimensional dipole, confirming again the perturbation model of Benisty.

For higher refractive index contrasts, the relationship ceases to be quadratic, and the losses level out. Therefore when one is working in the region of low to medium index contrast, like a $GaAs - Al_xGa_{1-x}As$ heterostructure, this simple model suggests it might be beneficial to keep the refractive index contrast as low as possible, for a small increase in the refractive index of the waveguide cladding might yield a considerable drop in losses. For high refractive index contrasts (Semiconductor membranes in air, or Silicon-on-insulator,...) a change in re-

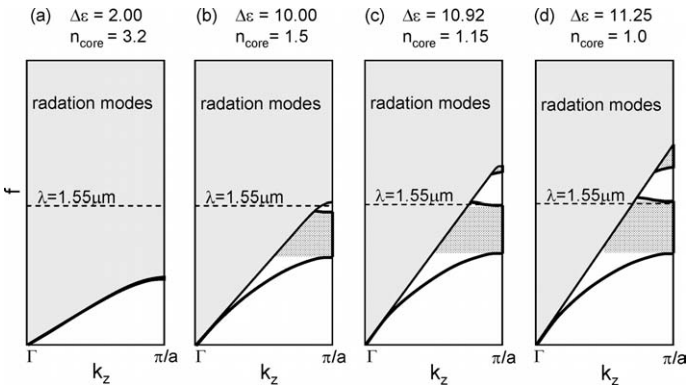

Figure 4.5: Band diagrams of 1-D photonic crystal slabs. The radiation modes are shaded, and the photonic bandgap is coarsely shaded. Left to right: (a) $\Delta\epsilon = 2.0$, (b) $\Delta\epsilon = 10.0$, (c) $\Delta\epsilon = 10.92$, (d) $\Delta\epsilon = 11.25$. The operating wavelength $\lambda = 1.55\mu m$ is also indicated.

fractive index of the cladding is far less important from the viewpoint of out-of-plane scattering losses.

4.2.3 Multiple Air Slots

While the simulation of a single air slot gives us some insights in the effect of the vertical index contrast on out-of-plane scattering, it totally neglects the fact that photonic crystals are periodic structures.

As we have seen in chapter 2, a 2-D periodic structure can have totally lossless Bloch modes, because the mode does not couple to the radiation continuum [39, 40]. In the approximation of the holes as scattering dipoles, this is possible when the scattering losses of each of the holes interfere destructively.

Such a guided lossless Bloch mode can of course also be supported by a 1-D periodic structure as we use here. We simulated the structure in the right part of figure 4.3, with a period a of $0.55\mu m$ and an air slot width d_a of $0.28\mu m$. These parameters are chosen such that at a wavelength of $1.55\mu m$ the periodic structure exhibits all possible scattering regimes when varying the refractive index contrast between core and cladding of the slab waveguide.

Figure 4.5 shows the band structure of the infinitely 1-D periodic structure for different values of the index contrasts $\Delta\epsilon$ between core and cladding layers. The shaded region bounded by the light line indicates the continuum of the radiation modes supported by the top

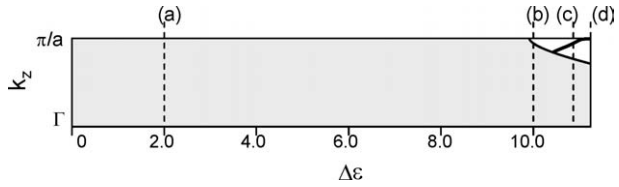

Figure 4.6: Modes of a 1-D photonic crystal slab at $1.55\mu m$ as a function of index contrast $\Delta\epsilon$. The situations of the four band diagrams from figure 4.5 are also indicated.

and bottom cladding. Figure 4.6 shows the propagation constants of the Bloch modes at a wavelength of $1.55\mu m$ as a function of refractive index contrast. We see that for high contrasts the structure quickly changes between four distinct regimes. In figure 4.5a ($\Delta\epsilon = 2.0$, typical for a $GaAs - Al_xGa_{1-x}As$ heterostructure) we are obviously working outside of a bandgap and above the light line. Therefore, we expect out-of-plane scattering losses because there is coupling to the radiation modes. As the index contrast increases, the perturbation by the air slots is larger, and we see that the losses increase with increasing refractive index contrast. The losses are plotted in figure 4.7 as a function of the refractive index contrast $\Delta\epsilon$. When working outside of the bandgap, the characteristics of a finite periodic structure are a non-monotonous function of the number of periods N. To eliminate these fluctuations, we took the average of the losses of the periodic structures with 30 to 50 periods. At an index contrast of $\Delta\epsilon = 10.00$ (figure 4.5b) the infinite periodic structure has a region below the light line for high values of k_z. However, there is no guided mode present in this region. At the slightly higher index contrast of $\Delta\epsilon = 10.92$ (figure 4.5c) there is a guided mode below the light line at the wavelength of $1.55\mu m$. For an infinite periodic structure, this mode would be truly lossless, as it wouldn't couple to the radiation modes. This is only true for an infinite structure, or in close approximation for the bulk of a finite periodic structure. We see in figure 4.7 that losses have dropped sharply for this structure, and the transmission of the structure has increased dramatically. In this regime, most losses are due to the transition between the homogeneous slab waveguide and the periodic structure. For even higher refractive indices (figure 4.5d) there are no guided modes either below or above the light line, as there is a bandgap at the operating wavelength. Most of the incident light is now reflected in the first few periods and can't penetrate the structure. Note that the bandgap is ac-

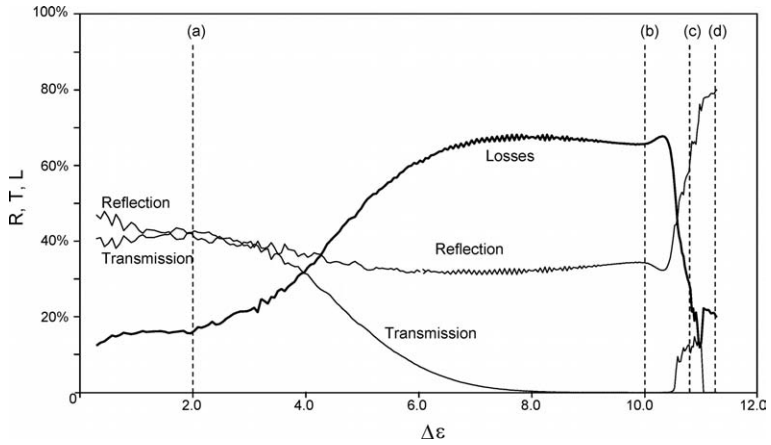

Figure 4.7: Reflection, transmission and out-of-plane scattering loss for a slab waveguide with multiple air slots as a function of the vertical refractive index contrast $\Delta\epsilon$. The values are averaged over a structure with 30 to 50 periods to average out oscillations related to the finite number of periods. The four situations outlined in figure 4.5 are indicated.

tually a Bloch mode that is exponentially attenuated in the propagation direction. As this mode is below the light line, there is no coupling to the radiation continuum. Therefore losses are also quite low in this regime, and can again be credited to transition losses at the boundary of the periodic structure.

In figure 4.8 we plotted the losses of these four structures as a function of the number of periods N. For each set of data we added a broad moving average. As expected, for low index contrast ($\Delta\epsilon = 2.0$) we find low losses per slot, but the total loss increases as more slots are added. For a higher index contrast of $\Delta\epsilon = 10.00$ losses per hole are much higher, and we need only a few holes for all the light to be lost (the remaining light is reflected). If there is a guided Bloch mode at the operating wavelength, the only losses we observe are coupling losses from the guided slab mode to the guided Bloch mode and back, so as expected the losses remain constant regardless of the number of periods, as does the transmission. This effect occurs already after a small number of periods. If we operate within the PBG of the structure, most incident light is reflected back into the waveguide within the first few periods, so losses remain constant as we add more periods. However, transmission quickly drops to zero in this regime as we increase the number of periods.

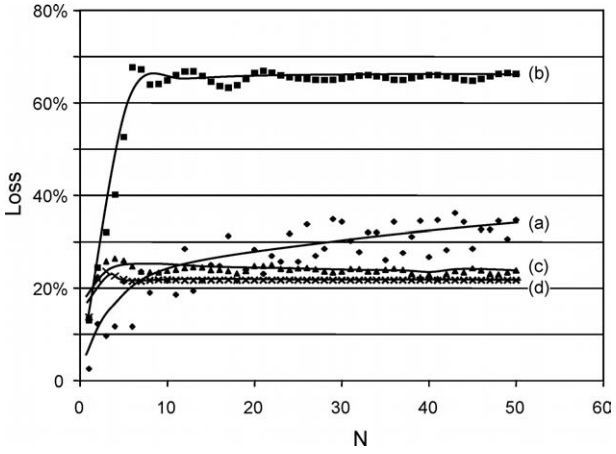

Figure 4.8: Out-of-plane scattering losses as a function of number of periods N in the 1-D photonic crystal slab for the four values of $\Delta\epsilon$ from figure 4.5. A moving average is also plotted for each case. For low contrast ($\Delta\epsilon = 2.0$), losses are low but increase slowly with each added period. For higher contrasts ($\Delta\epsilon = 10.0$) the losses climb rapidly and reach a maximum after only a small number of periods. Only the reflected power is not lost. As a guided Bloch mode is excited around $\Delta\epsilon = 10.92$, the losses remain constant regardless of the number of periods, and the only loss is due to the transition between the waveguide and the periodic structure. When we operate within the PBG $\Delta\epsilon = 11.25$, only a few periods are needed to reflect most of the power back into the waveguide and the losses remain constant with the number of periods N.

Note that the period a and the slot widthd_a were chosen such that we went through the four regimes by changing the index contrast. For other parameters, the conclusions are still valid, but the presence of regime c (guided Bloch mode) and/or d (PBG below the light line) is not guaranteed.

4.2.4 Conclusions for 2-D Photonic Crystal Slabs

Losses due to out-of-plane scattering put a limitation to the usefulness of photonic crystal slab waveguides. Designing the optimum layer structure to keep these losses low is therefore an important step in the design of these structures. We have shown that if we treat the losses of each hole independently, losses increase with the square of $\Delta\epsilon$ for low refractive index contrasts, as predicted by Benisty in 1999 [36, 37]. For higher refractive index contrasts however, this trend levels out. Because this single-slot simulation is only a coarse approximation, we also studied the out-of-plane scattering losses of a periodic structure with multiple air slots. We can distinguish four different regimes for such a structure. For low refractive index contrasts, the behavior is similar as with the single air slot. Losses increase with increasing refractive index contrasts. For high contrasts losses will be very high unless one can excite a lossless Bloch mode below the light line [40], resulting in a drop in out-of-plane scattering losses and a higher transmission of the structure. If one operates in a bandgap of the photonic crystal below the light line, losses will be low because light cannot penetrate the crystal very deeply.

From the perspective of out-of-plane scattering, two regimes are favorable to reduce the intrinsic out-of-plane scattering losses, depending on the application. The principle is illustrated in figure 4.9. For applications with many defects (bends, splitters, ...) close to each other, a low refractive index contrast between core and cladding is well suited, because the holes scatter similar amounts of light in the bulk of the periodic structure and where the periodicity is disturbed at the boundaries of the structure. This suggests that in a 2-D photonic crystal slab waveguide, intentional defects like bends and splitters will not behave significantly different from the straight waveguides. This makes it possible to approximate the air holes as made of a lossy material, and assigning a complex refractive index. Further designs can then be made in 2-D, which significantly speeds up the calculations. Benisty has shown that this technique can be used to obtain very good qualitative and

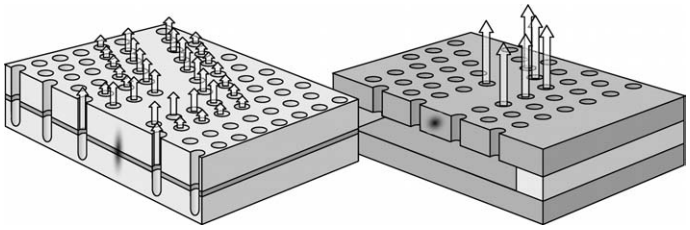

Figure 4.9: Regimes with low intrinsic out-of-plane scattering. Left: low vertical index contrast. Light is scattered little by little as is propagates, both in the waveguide and in defects. Right: High vertical index contrast. When a guided Bloch mode is excited, no losses occur in the straight waveguide, but a breach in periodicity can trigger large out-of-plane scattering.

even quantitative results, and theory agrees reasonably well with experimental observations [37].

However, if one has few and/or widely separated defects in the waveguide, it is preferable to couple light into a lossless Bloch mode. This requires a high refractive index contrast in the vertical direction. The guided Bloch mode in a 2-D photonic crystal slab will not scatter light as long as the periodicity of the waveguide is not disturbed. However, when a defect is introduced, this can cause a large amount of scattering. Therefore, thorough optimisations of these defects is necessary, and the design should be done in 3-D.

In this section, we calculated the losses of a 1-D periodic structure with air slots that extend infinitely in the vertical direction. This is an approximation of slots that are etched very deep into the bottom cladding of the slab waveguide, where the guided mode is sufficiently confined to the waveguide core. Additional losses are to be expected when the slots aren't etched deep enough or with insufficient confinement, e.g. when the index contrast is too low, and the waveguide mode is very broad. For realistic structures in the low contrast regime, one will therefore have to balance between keeping a low index contrast (material choice) and the depth of the air slots (etching facilities) to keep the losses as low as possible. In the high index contrast regime, one can target a regime where the losses in the bulk of the periodic structure can theoretically be reduced to zero. However, applications using these structures will always require defects that will introduce losses in the otherwise lossless structure [38, 41].

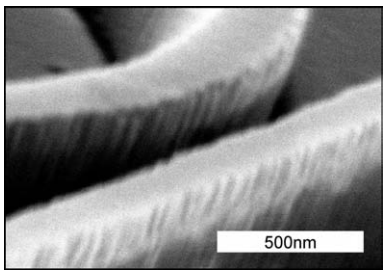

Figure 4.10: Etch roughness in photonic wires. This is an example of a deeply etched SOI waveguide with significant etch roughness. The roughness consists of vertical grooves in the sidewalls of the waveguide.

4.3 Roughness in Photonic Wires

While some photonic crystals have intrinsic scattering losses, photonic wires are in principle lossless waveguides as long as the cross section doesn't change along the propagation direction. However, in real fabricated structures, photonic wires will have some irregularities due to the imperfect fabrication process. For instance, with dry etching techniques there is always a certain amount of sidewall roughness. An extreme example of such a rough sidewall is given in figure 4.10 for a deeply etched SOI waveguide.

While light in a photonic crystal slab can sometimes be 'recovered' because it can't penetrate the photonic crystal cladding, scattered light from a photonic wire can freely propagate into the cladding in all directions. We have studied the effect of the refractive index contrast on the scattering losses in wires. This topic has already been studied in literature, and it is agreed that losses increase dramatically with refractive index contrast, but sources disagree on the exact relationship. It has been reported that the losses are proportional to either Δn^3 [42, 43], $\Delta\epsilon$ [44, 45] or $\Delta\epsilon^2$ [46], with $\Delta\epsilon = n_{core}^2 - n_{clad}^2$.

Another important question with respect to waveguide losses is the choice of polarisation. With a specific type of roughness (e.g. roughness due to etching) is it better to use the TE-like polarisation in the waveguide or the TM-like polarisation in order to minimise the scattering losses due to roughness?

To compare the scattering at roughnesses in photonic wires for different index contrasts and both polarisations, we model the scattering center (i.e. an irregularity on the sidewall) by a radiating dipole. Figure 4.11 shows the simulation setup for the TE and the TM mode of the

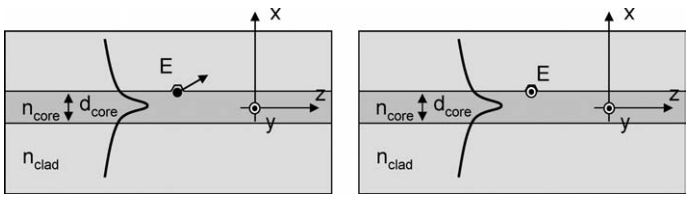

Figure 4.11: Scattering at sidewall roughness in photonic wires (top view). The roughness is treated as a radiating dipole on the sidewall of the wire. Left: For the TE mode in the wire (TM simulation). Right: For the TM mode of the wire (TE simulation).

wire. Ideally this problem is treated in 3-D, especially when the effect of the vertical index contrast (i.e. the choice of layer structure) is studied. However, we limit ourselves to a 2-D comparison of the TE and TM mode of the photonic wire for different lateral index contrast.

The simulated geometry represents a top view from the photonic wire waveguide, because we want to study the effect of etch roughness. This type of roughness typically consists of vertical grooves in the waveguide, as illustrated in figure 4.10. Because we use a top view, the TE mode of the wire will be simulated by a TM simulation, while the TM mode of the wire will be simulated with a TE simulation. The simulations were done using the eigenmode expansion tool CAMFR.

We model the light scattered in the sidewall by a radiating dipole. For a dipole with field strength E_0 on the sidewall, the radiated power P_0 is

$$P_0 = \frac{E_0^2}{Z_{rad}} \tag{4.1}$$

with Z_{rad} the radiative impedance of the environment, which of course differs depending on the orientation of the radiating dipole. The radiating dipole will be excited by the guided mode of the waveguide, so we expect that the field E_0 is proportional to the field strength E on the sidewall of the waveguide, and has the same direction. The dipole will therefore radiate a power P proportional to the square of the local field E:

$$P = \eta^2 \frac{E^2}{Z_{rad}}. \tag{4.2}$$

Depending whether we look at the quasi-TE or TM mode of the waveguide, E is represented by the in-plane or the vertical electric field respectively, as illustrated in figure 4.11. For high index contrasts and narrow wires, the electric field of the TE mode on the interface has a

significant component in the z direction, therefore

$$\begin{aligned} \text{TE mode:} \quad & E \approx \sqrt{E_x^2 + E_z^2} \\ \text{TM mode:} \quad & E \approx E_y \end{aligned} \tag{4.3}$$

As we have normalised the guided modes to unity power, the respective fields in both polarisations are normalised with respect to each other. Also, where the electric field has a discontinuity at the interface (as E_x for the TM mode), we have always used the largest value of E at the discontinuity.

The factor η^2 describes the influence of the roughness geometry and the horizontal index contrast $\Delta\epsilon_h$ on the sidewall. We also approximate the radiative impedance of the perturbed structure by the impedance of the smooth waveguide.

We calculate a rough estimate for η^2 by scattering a plane wave off an irregularity on a material-air interface. We do this for different shapes of irregularities and $\Delta\epsilon_h$. Figure 4.12 shows η^2 as a function of $\Delta\epsilon_h$ for different irregularities. When the irregularity is parallel to the incident E-field (as for the TM mode of the waveguide) η^2 behaves very closely like

$$\eta^2 = \gamma_{\parallel} \left(\Delta\epsilon_h\right)^2 , \tag{4.4}$$

with $\gamma_{\parallel}$ describing the influence of the shape and size of the irregularities in the case that the incident E field is parallel with the orientation of the irregularity on the interface. This relationship with $\Delta\epsilon_h$ is also confirmed in literature for orthogonal incidence on an interface. For oblique incidence, the absolute scattering is lower, but the relationship with $\Delta\epsilon_h$ remains largely unchanged, especially for high index contrasts [46].

When the E-field is perpendicular to the axis of the irregularity on the interface (like for the TE mode of the waveguide), we find a similar relationship. We can see in the bottom part of figure 4.12 that η^2 is again more or less proportional to the square of $\Delta\epsilon_h$, especially for the smaller irregularities. Because this is a better approximation of sidewall roughness, we can also write

$$\eta^2 = \gamma_{\perp} \left(\Delta\epsilon_h\right)^2 . \tag{4.5}$$

We can now compare the scattering strength of a dipole in the case of the TM polarisation and the TE polarisation in the waveguide by taking the ratio between $\gamma_{\parallel}$ (for the TM polarisation) and $\gamma_{\perp}$ (for the TE polarisation). This ratio is plotted in figure 4.13. Then we take a

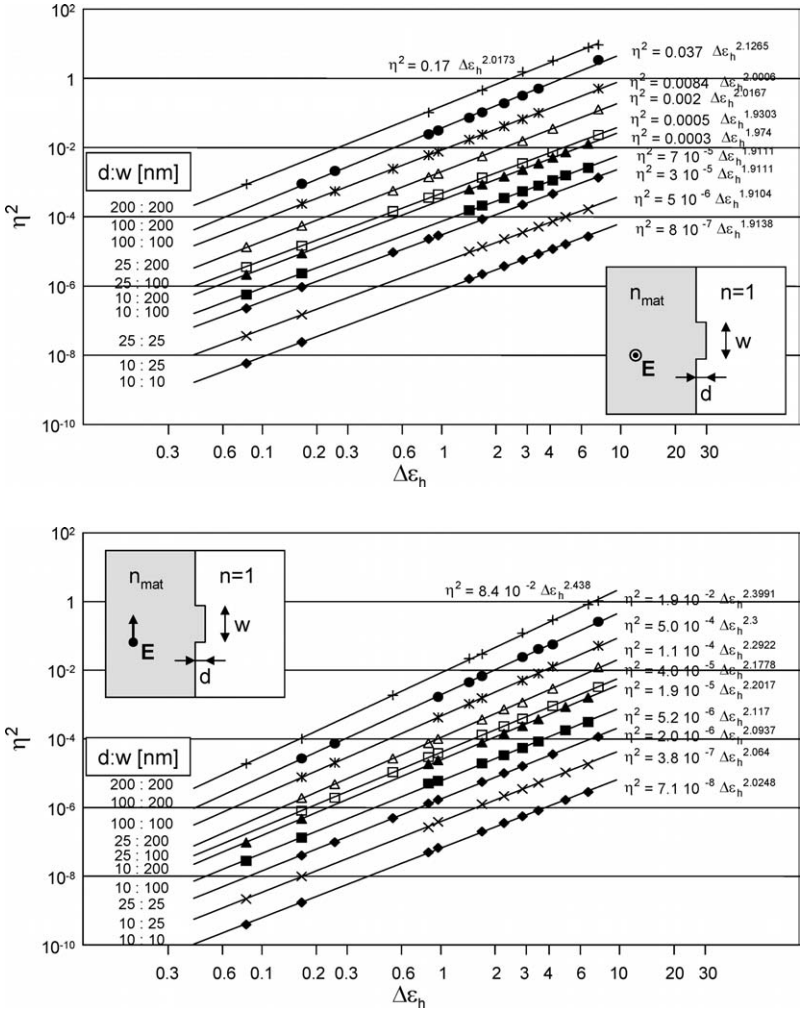

Figure 4.12: Strength of a 2-D dipole excited by a plane wave incident on a 2-D irregularity on a material-air interface, plotted as a function of the horizontal index contrast $\Delta\epsilon_h$ for different irregularities. Top: the plane wave is polarised with the E-field parallel to the axis of the irregularity. Bottom: the plane wave is polarised with the E-field perpendicular to the axis of the irregularity

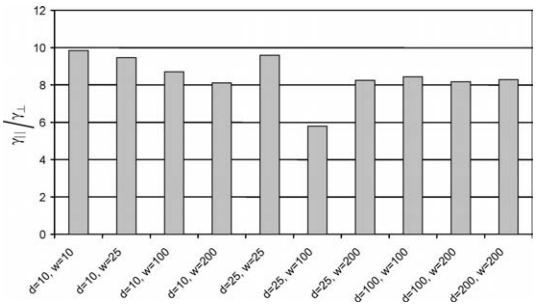

Figure 4.13: Ratio between $\gamma_{\parallel}$ and $\gamma_{\perp}$ for scattering at identical irregularities. The ratio is calculated by fitting the best fit for $\eta^2 = \gamma \left(\Delta\epsilon_h\right)^2$ to the data of both polarisations, and then taking the ratio of the resulting coefficient γ.

weighed average of $\gamma_{\parallel}/\gamma_{\perp}$ over the different irregularities, giving the small irregularities more weight as they more closely represent sidewall roughness. We find that

$$\frac{\gamma_{\parallel}}{\gamma_{\perp}} \approx 9. \tag{4.6}$$

The light is scattered much more when the electric field is parallel to the axis of the irregularity compared to the case where E is perpendicular to the irregularity. This discrepancy between polarisations is also demonstrated in literature with the scattering off a cylinder partially embedded in a dielectric [89]. It can be explained by the discontinuities of the E field perpendicular to an interface. The scattering cross section of the 2-D scatterer will be larger when E is parallel to the scatterer.

We can now write the power radiated by the dipole for both the TE and TM mode. For the TE mode we will have to use $\gamma_{\perp}$ as the E field is largely oriented along the x axis, while the sidewall roughness is oriented along the y axis. For the TM mode, we will use $\gamma_{\parallel}$, as E is strongest in the the y direction.

$$\text{TE mode:} \quad P_{TE} = \gamma_{\perp} \left(\Delta\epsilon_h\right)^2 \frac{E_x^2}{Z_{rad,TE}} \tag{4.7}$$

$$\text{TM mode:} \quad P_{TM} = \gamma_{\parallel} \left(\Delta\epsilon_h\right)^2 \frac{E_y^2}{Z_{rad,TM}}. \tag{4.8}$$

Because the dipole radiates in all directions, part of the scattered light will be recaptured by the guided mode of the waveguide travelling in the forward direction. Using CAMFR, we can easily extract the amount of power from the dipole that is coupled to the guided mode.

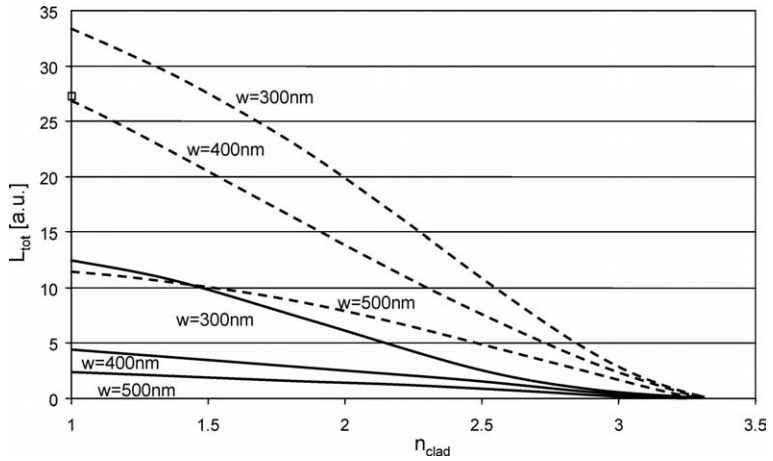

Figure 4.14: Measure for the amount of scattering at sidewall roughness in photonic wires. The measure L_{tot} for the scattering loss is plotted as a function of the lateral index contrast $\Delta\epsilon_h$ of the waveguide. The results for 3 wire widths and both TE (full line) and TM polarisations (dashed line) is shown.

We call L the fraction of the power that is not recovered by this guided mode. This fraction will be different for TE or TM. The total lost power L_{tot} can therefore be written as:

$$\text{TE mode:} \quad L_{tot,TE} = \gamma_{\perp} \left(\Delta\epsilon_h\right)^2 L_{TE} \frac{E_x^2}{Z_{rad,TE}}, \tag{4.9}$$

$$\text{TM mode:} \quad L_{tot,TM} = \gamma_{\parallel} \left(\Delta\epsilon_h\right)^2 L_{TM} \frac{E_y^2}{Z_{rad,TM}}. \tag{4.10}$$

Both relations are plotted in figure 4.14 for 3 different widths of wires and for both polarisations as a function of the cladding index n_{clad}. The units are arbitrary, as the actual amount of scattering is strongly dependent on the geometry of the roughness, but all curves are drawn on the same scale. For low index contrasts, the losses are low, but then the waveguides can hardly be classified as wires. For higher contrasts, we clearly see higher losses as the wire width decreases, but this is hardly surprising as the field strength on the sidewall of the wire is higher for narrow wires. We also see that the scattering losses for the TM mode are significantly higher than for the TE mode.

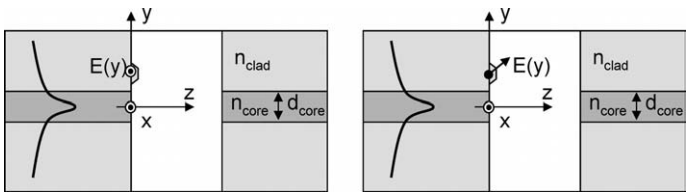

Figure 4.15: Out-of-plane scattering in a single air slot (side view). The roughness is treated as a radiating dipole. Left: For TE polarisation. Right: For TM polarisation.

4.4 Roughness in Photonic Crystals

In photonic crystal slab waveguides, scattering at sidewall roughness is also an important source of roughness. In addition, sidewall roughness can disturb the periodic nature of the photonic crystal. Just as with intrinsic scattering losses in photonic crystal slabs, we have studied the effect of sidewall roughness with respect to the vertical index contrast.

Intuitively, it is not easy to favour either high index contrast or low index contrast. With high index contrast, the guided mode contains strong out-of-plane components, but because of the high index contrast, the waveguide can also recapture scattered light at large angles. Also, the intensity of light on the surface is higher, as the light is more confined in the core. In the other case, the light behaves more like a plane wave, and contains little out-of-plane components. However, it is stretched out wider and therefore feels more surface area.

4.4.1 Scattering Model

To compare the situation of high and low vertical index contrast, we modelled the sidewall irregularities as radiating dipoles excited by the incident guided light. Again, we are limited to the 2-D approximation of air slots in a slab waveguide. Because we are mainly interested in photonic crystal slabs with etched holes, we used the TE polarisation, and light at a wavelength of $1550nm$. The geometry of the simulated structure is shown in figure 4.15. The CAMFR tool is very well suited for this type of structure, as it can model the radiation of dipole sources, together with the use of PML boundary conditions.

To compare slabs with different vertical index contrast, we used a three-layer symmetric slab waveguide with a v-number of 2.79 to guar-

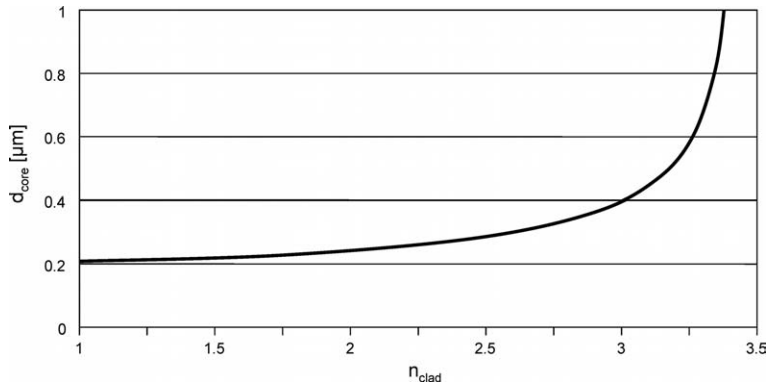

Figure 4.16: Waveguide core width d_{core} as a function of cladding index n_{clad}, for the geometry in figure 4.15. The core width is changed in order to keep the waveguide's v-number constant for a constant v-number

antee the single-mode behaviour. The v-number is related to the refractive index contrast and the thickness of the waveguide:

$$v = k_0 d_{core} \sqrt{n_{core}^2 - n_{clad}^2} \tag{4.11}$$

To keep the v-number constant, the core thickness d_{core} increases with the cladding index n_{clad}, as illustrated in figure 4.16. For compatibility with both $GaAs/AlGaAs$ and Silicon-based structures, the index of the core $n_{core} = 3.45$ was chosen.

When guided light hits the sidewalls, it will be scattered by irregularities. Because these irregularities are not necessarily the same in each slot or hole, we assumed that the total power lost L_{tot} due to roughness at the sidewall is proportional to the power loss of an irregularity averaged over any position y along the interface:

$$L_{tot} \sim \int_y P(y)\, L(y)\, dy \tag{4.12}$$

with $P(y)$ the scattered power of the dipole at position y, and $L(y)$ the fraction of that power that is not recaptured by the waveguide, and therefore lost. Because the dipole is excited by the incident field, $P(y)$ can be written as a function of the local field $E(y)$:

$$P(y) = \eta^2(y) \frac{E^2(y)}{2Z_{rad}(y)} \tag{4.13}$$

with $\eta(y)$ describing the effect of the roughness geometry on the interface, and $Z_{rad}(y)$ the radiative impedance of the environment. $E(y)$ is the electric field of the guided slab mode in position y.

4.4.2 Simulations

We use 3 simulations to calculate the different factors in this model. First, we determine $E(y)$ and $Z_{rad}(y)$ by simulating the response of a dipole current source on the interface. Then we approximate η^2 by simulating the dipole from an irregularity on an interface excited by a plane wave. Finally, we calculate the local E-field from the slab mode.

In the TE polarisation, a 2-D dipole radiates in all directions. In the TM polarisation, the radiation of a 2-D dipole has two lobes in the horizontal direction. We studied both TE and TM polarisation. We will discuss the case for TE in detail, but the procedure for TM is similar, although the numerical simulations with CAMFR proved to be more difficult.

The slab core will recapture a fraction of the radiated power. In a photonic crystal the forward and backward propagating light is coupled, so the recaptured light in both forward and backward direction is considered not lost. The fraction of lost light $L(y)$ can be expressed as

$$L(y) = \frac{P_0(y) - R(y) - T(y)}{P_0(y)} \tag{4.14}$$

with $R(y)$ and $T(y)$ the power in the forward and backward propagating guided mode, and $P_0(y)$ the dipole power. We calculate this by putting a dipole current source with a known current I_0 on the interface at position y. The power $P_0(y)$ emitted by this source is

$$P_0(y) = \frac{I_0 E_0(y)}{2} \tag{4.15}$$

where $E_0(y)$ is the electric field at the position of the source. As this is a 2-D simulation, $P_0(y)$ is measured in $[W/m]$. Likewise, the radiative impedance of the environment

$$Z_{rad} = \frac{E_0(y)}{I_0} \tag{4.16}$$

is expressed in $[\Omega/m]$. Figure 4.17 shows $L(y)$ for different values of n_{clad}. If the dipole is located near the slab core, a larger fraction of light is recovered. Also, a slab waveguide with high vertical index contrast can recapture more light.

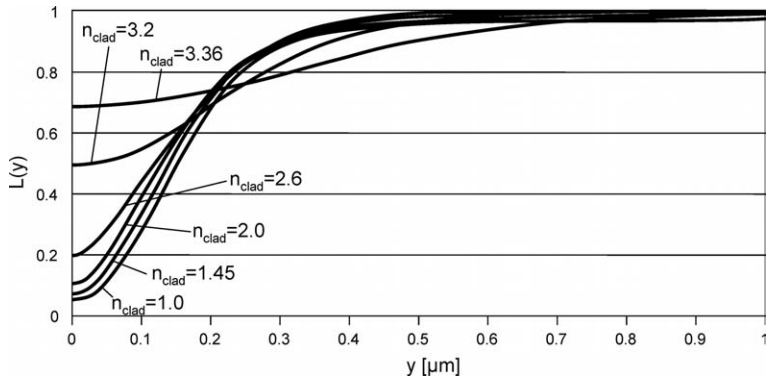

Figure 4.17: Loss of a radiating dipole (E in the x direction) on the sidewall of a slot as a function of position y. $L\,(y)$ is the fraction that cannot be recaptured by the waveguide.

We then calculate the excitation of the dipole η^2 by the incident field. It is obvious that the amount of scattering depends strongly on the type of irregularities and the horizontal index contrast $\Delta\epsilon_h = n_{mat}^2 - 1$ of the sidewall, with n_{mat} the index of the material at position y.

As with photonic wires, we approximate the effect of the roughness η^2 by simulating the incidence of a plane wave on an interface with an irregularity. As we can see in figure 4.12 on page 108 η^2 is proportional to the square of the difference in relative dielectric constant on the interface $\Delta\epsilon_h$:

$$\eta^2 = \gamma\,(\Delta\epsilon_h)^2\,. \tag{4.17}$$

The factor γ describes the influence of the shape and size of the irregularities. As the most common type of roughness consists of vertical grooved caused by dry etching, we will use $\gamma_\perp$ for the TE guided mode as the electric field is oriented perpendicular to the grooves. Likewise, we will use $\gamma_\parallel$ for the TM guided mode, which is typically an order of magnitude larger than $\gamma_\perp$.

4.4.3 Results

Knowing η^2 and approximating $Z_{rad}\,(y)$ by assuming that Z_{rad} of the environment doesn't differ much when a small irregularity is applied to the sidewall, we now calculate the excitation $P\,(y)$ of the dipole at position y:

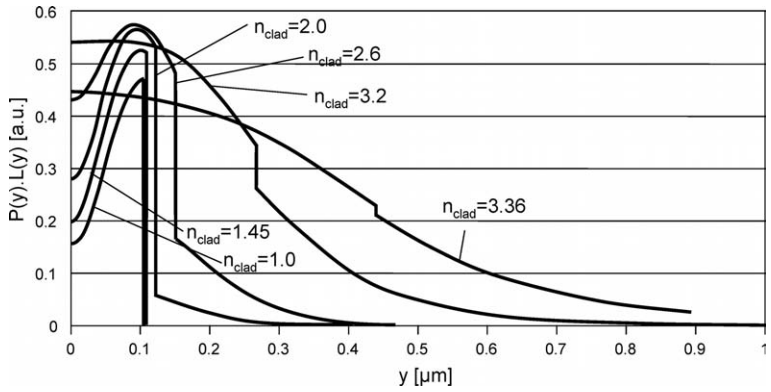

Figure 4.18: Lost power power of a dipole at position y on the sidewall excited by the TE polarised slab waveguide mode incident on a single air slot. The discontinuities occur at the core-cladding interface.

$$P(y) \approx \gamma \left(\Delta\epsilon_h(y)\right)^2 \frac{E^2(y)}{2Z_{rad}}. \tag{4.18}$$

Note that $\Delta\epsilon_h$ is dependent on the position y, as the index contrast of the material-air interface is higher in the slab core than in the slab cladding. Therefore, as the core index n_{core} is fixed, a low vertical index contrast implies a high value of $\Delta\epsilon_h$ for the cladding. Figure 4.18 shows $P(y)$ for a different n_{clad}, with $E(y)$ based on the guided slab mode. Overall, the dipole excitation is much stronger for low vertical index contrasts because the mode profile is much broader, and $\Delta\epsilon_h$ in the cladding is much larger.

We can now combine the results to estimate the average losses due to roughness. Figure 4.19 shows L_{tot} as a function of n_{clad}. There is a strong increase of losses for high values of n_{clad}, i.e. low vertical index contrast. We therefore conclude that high vertical index contrast (like SOI or semiconductor membranes) performs better with respect to losses at irregularities than low out-of-plane contrast (like $GaAs/AlGaAs$ waveguides).

We did the same simulation for the TM polarisation. The result is also plotted in figure 4.19. When we compare both polarisations in the waveguide, we find that they have similar behaviour as a function of refractive index contrast. However, just as with photonic wires, the TM polarisation is more prone to scattering.

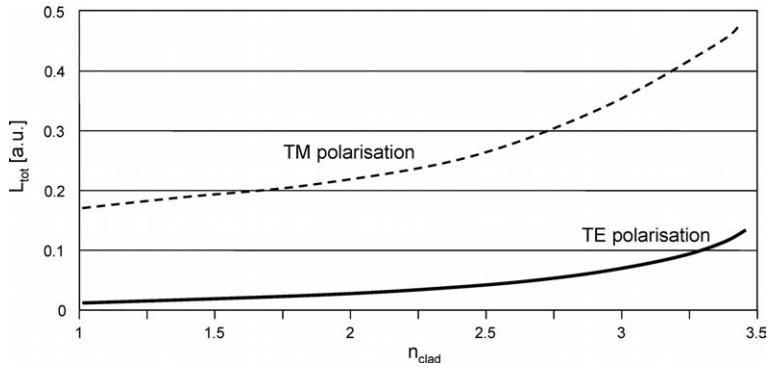

Figure 4.19: Average power lost due to scattering at irregularities on the sidewall of a single air slot as a function of the slab cladding index n_{clad}. The TE polarisation is plotted in a full line, while the result for TM are plotted in a dashed line.

4.4.4 Conclusion

Intrinsic out-of-plane scattering favours two regimes with respect to propagation losses in photonic crystal slabs. However, when we consider real structures with sidewall roughness due to imperfect fabrication, additional scattering losses occur similar to those in photonic wires. On the sidewall of the etched holes, light is scattered out-of-plane. This scattering increases strongly with the in-plane refractive index contrast on the vertical interface. Because low vertical index contrast (i.e. a high-index cladding layer) implies a strong in-plane index contrast, the scattering losses due to roughness are several times larger then for high vertical index contrast materials such as membranes or SOI. This might partially explain why SOI photonic crystal waveguides still outperform photonic crystal waveguides in III-V heterostructures [47, 48]. We also found that TE polarisation performs much better with respect to scattering at sidewall roughness than the TM polarisation. Because the TE polarisation also has a PBG in a photonic crystal slab with a triangular lattice of holes, this is the preferred polarisation.

4.5 Substrate Leakage

Up till now we have always studied nanophotonic waveguides fabricated from a three-layer slab waveguide with an infinite top cladding, a guiding core and an infinite bottom cladding. In real life, the bottom

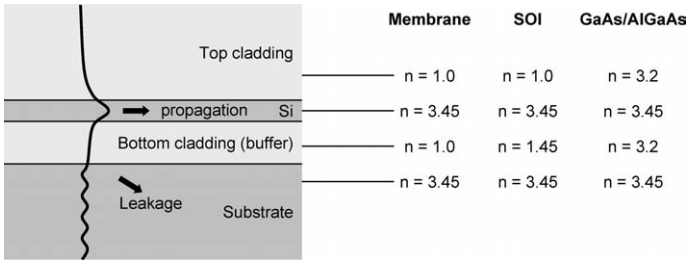

Figure 4.20: Substrate leakage losses. If the substrate has the same or higher refractive index as the waveguide core, light can leak through the bottom cladding towards the substrate.

cladding and sometimes the top cladding have a finite thickness. Typically, the slab waveguide is deposited or grown onto a semiconductor substrate. In the case of SOI or $GaAs/AlGaAs$, the substrate consists of pure Silicon or $GaAs$, which has the same refractive index as the waveguide core, and cannot act as a bottom cladding. Therefore, the bottom cladding of the waveguide consists of one or more layers of a lower-index material. However, when this layer has a finite thickness, light can leak from the waveguide core towards the high-index substrate, because the exponential decaying tail of the guided waveguide mode can 'feel' the high-index substrate. This is illustrated in figure 4.20. Because the intensity of the waveguide mode decreases exponentially in the waveguide cladding, we expect that the amount of substrate leakage behaves similarly. Therefore, in order to avoid substrate leakage, a sufficiently thick buffer layer is needed.

4.5.1 Slab Waveguide

Using CAMFR, we have simulated the substrate leakage for silicon-on-insulator material as a function of core thickness and oxide cladding thickness. Because CAMFR is a complex eigenmode solver, it calculates both the real part and the imaginary part of the propagation constant of the waveguide modes. The imaginary part corresponds to either gain or loss along the propagation axis. If we use only lossless material, the only loss mechanism will be leakage to the substrate. Because this is an open structure, perfectly matched layers are needed to model the radiation modes in the substrate.

We simulated the substrate leakage for a Silicon-on-insulator slab waveguide, as illustrated in figure 4.21 as a function of oxide thick-

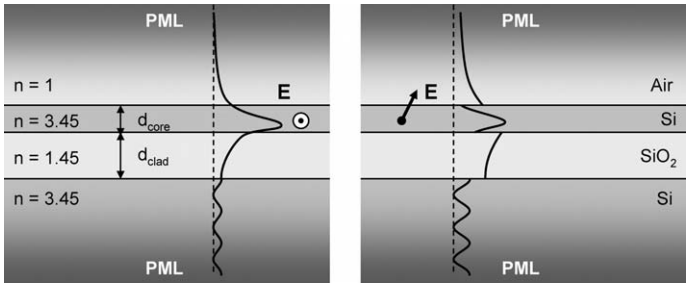

Figure 4.21: Simulation setup for substrate losses in a slab waveguide. Left: TE mode, Right: TM mode. Because of discontinuities at the core-cladding interface, the TM mode reached deeper into the cladding.

ness d_{clad}. As shown in figure 4.21, the TM mode of the waveguide reaches deeper into the buffer layer because of the discontinuities at the layer interfaces. Therefore, TM polarised light is more prone to substrate leakage. The simulation results are plotted in figure 4.22. A number of trends are immediately obvious. Firstly, the TM mode has indeed much higher substrate leakage losses than the TE mode. Also, the substrate leakage decreases exponentially as a function of cladding thickness d_{clad}. The results are plotted for a number of core thicknesses d_{core}. As the core is thicker, the mode is better confined and the substrate leakage is also reduced. However, for a certain core thickness ($d_{core} > 255nm$, the core ceases to be single-mode (for each polarisation).

Because the waveguide mode changes as a function of wavelength, the substrate leakage will also change. Figure 4.23 plots the propagation loss due to substrate leakage as a function of wavelength between $1.4\mu m$ and $1.7\mu m$ for a number of buffer thicknesses. We can see that the substrate leakage increases with wavelength. As the core is less thick compared to the wavelength, the mode is less confined.

For the fabrication of our structures, we have ultimately chosen a SOI core thickness of $220nm$. In that case, we can conclude that a buffer layer of $1\mu m$ is sufficient to eliminate most of the substrate leakage for the TE polarisation. However, for TM a buffer layer of at least $2\mu m$ is needed.

4.5.2 Photonic Wire

The substrate leakage losses of a slab waveguide give us a lower limit of the propagation losses. In real life, we will use waveguides of a

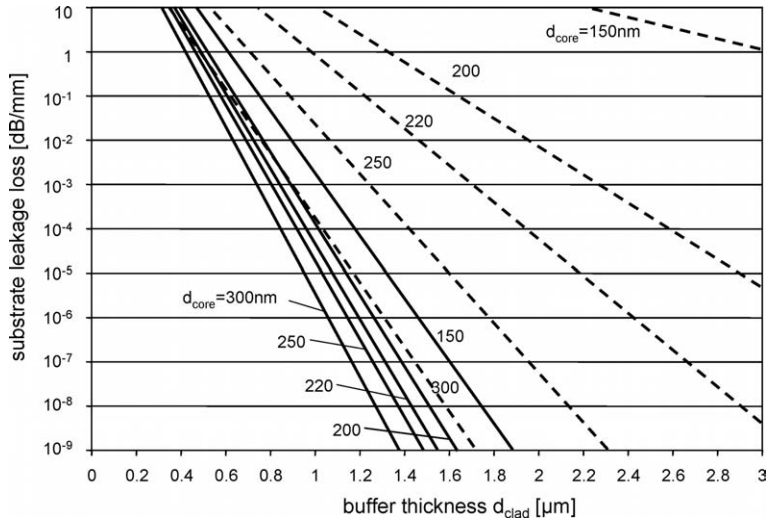

Figure 4.22: Substrate leakage losses for a slab waveguide as a function of buffer thickness d_{clad}. The result is also plotted for different core thicknesses d_{core}. TE polarisation is marked in full line, while the results for TM are marked in dashed lines.

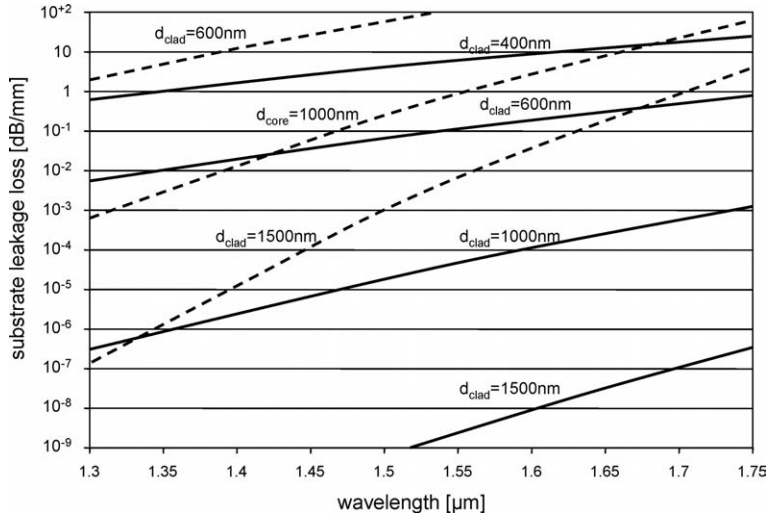

Figure 4.23: Substrate leakage losses for a slab waveguide as a function wavelength for different buffer thicknesses d_{clad}. The core thickness d_{core} is $220nm$. TE polarisation is marked in full line, while the results for TM are marked in dashed lines.

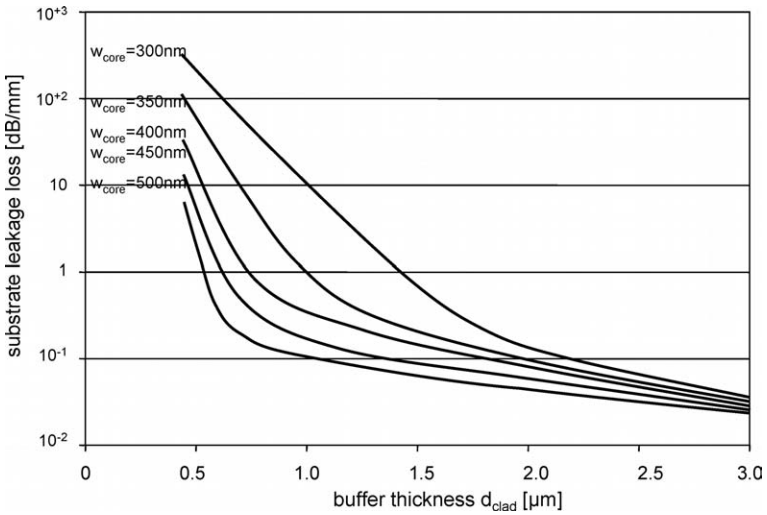

Figure 4.24: Substrate leakage for photonic wires as a function of cladding thickness d_{clad} with various core widths w_{core}.

finite width. For broad, multimode waveguides, the loss figures will be similar to the slab waveguides, as the ground mode closely resembles the slab mode.

However, in photonic wires, the core width is of the same order as the core thickness. As already discussed in chapter 3, because the core is now flanked by low-index material, the vertical confinement is reduced, and the mode will reach deeper into the bottom cladding layer. This effect is illustrated in figure 3.21 on page 78 in chapter 3. This will increase the losses due to substrate leakage.

The losses for the guided TE mode are plotted in figure 4.24 as a function of cladding thickness d_{clad} for various wire widths. We see that the losses are dramatically larger than those calculated for a slab waveguide. For a wire with a width $w = 500nm$, substrate leakage losses are already larger than $0.1dB/mm$,

If we increase the width of the core, the photonic wire will increasingly resemble a slab waveguide, and the substrate leakage losses will drop. This is illustrated in figure 4.25 for a cladding thickness $d_{clad} = 1\mu m$. The value for the slab waveguide is also illustrated. While broad photonic 'wires' have low substrate losses, they are no longer single-mode, and therefore, are not suitable as nanophotonic waveguides.

For photonic crystal slab waveguides, the effect is similar to that of photonic wires, as the mode is also strongly confined in the lateral

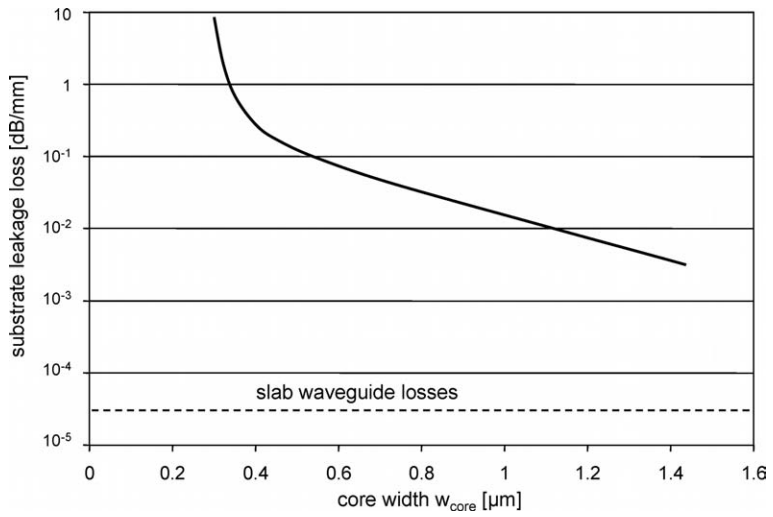

Figure 4.25: Substrate leakage for photonic wires as a function of core width w_{core} for a cladding thickness $d_{clad} = 1\mu m$.

direction. However, as the core layer is not completely removed on both sides of the waveguide core, the effect will probably be somewhat less. Because of the complexity of a photonic crystal waveguide, we were unable to simulate these structures rigorously.

4.6 Other Loss Mechanisms

We have discussed a number of loss mechanisms in photonic crystal slabs. Out-of-plane scattering has been identified as a prime loss mechanism, either intrinsic to the structure, or induced by sidewall roughness. We will now touch upon a number of additional loss mechanisms in photonic crystal slabs.

4.6.1 Sources of Scattering

Below the light line, a perfect photonic crystal waveguide can be lossless. When operating above the light line, all photonic crystal holes should be considered as sources of scattering. Even in a perfectly fabricated structure, there will be scattering losses. In real structures, fabrication is not perfect, and a number of other scattering mechanisms are introduced. Sidewall roughness is an important one, but a number of

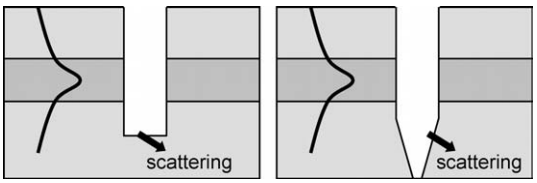

Figure 4.26: Additional loss mechanisms in photonic crystal slabs. Left: Scattering off the bottom of the etched holes. Right: sloped sidewalls in the bottom cladding.

other scattering sources are also possible. These are illustrated in figure 4.26.

When the etched holes are not etched deep enough, the light can scatter off the bottom of the holes. Also, in structures with a low vertical index contrast, the unetched cladding below the bottom of the hole can have an average refractive index that is higher than that of the perforated core. This will of course cause substrate leakage, especially since the mode extends deeply into the cladding for layers with low vertical index contrast.

A similar situation occurs when the sidewalls are sloped. This, again, is caused by imperfect etching techniques. When the etched holes get narrower, the average refractive index of the cladding increases and the mode is pulled downwards. Again, as the slab mode extends much deeper into the cladding for low vertical index contrast, the effect is more pronounced than with high vertical index contrast.

It is obvious that low vertical index contrast requires much better etching technology than membranes or SOI, for two reasons. In high-index-contrast structures, the core is often the top layer to be etched, while in low-index-contrast structures the core is deeply embedded in the top cladding. So deep etching is already needed just to penetrate the core. Additionally, the mode extends much deeper in the bottom cladding, requiring even more etch depth. Etch depths of more than $2\mu m$ are not uncommon in these structures, while for SOI $1\mu m$ already suffices.

4.6.2 TE/TM Interaction

When using a photonic crystal slab with a PBG for only one polarisation, the other polarisation is not confined by the PBG effect. In the case of a triangular lattice of holes in a high-index slab, like SOI, only the TE polarisation has a photonic bandgap. If we now create a waveguide by

removing or changing a row of holes, TE polarised light will be guided by the PBG. For the TM polarisation, only index-guided modes will be guided.

When the photonic crystal slab is symmetric in the vertical direction, TE and TM are perfectly orthogonal, and no interaction between the polarisations occurs. This is the case in Silicon membranes, where both the top and bottom cladding consist of air.

However, in Silicon-on-insulator, where the top cladding has a different refractive index from the bottom cladding, TE and TM modes will interact, forming modes of a hybrid nature [90]. If these hybrid modes are below the light line and outside the regions of both TE and TM lattice modes, they are still guided in the straight periodic waveguide (i.e. as long as the translational symmetry is unbroken) This has a number of undesirable effects. Firstly, the PBG-guided TE mode can couple to an index-guided TM mode. While both modes are guided in the straight waveguides, the TM mode is not as well confined in-plane in the case of bends and other defects. Also, the TM mode propagates at a different speed in the waveguide. Another effect is the possible occurrence of a so-called mini-stop-band (MSB). Because TE and TM modes interact, their dispersion relations will not cross, but the bands will be repelled, creating a mini-stop-band where no light can propagate. This is illustrated in figure 4.27.

The coupling between both polarisations can be reduced by increasing the vertical symmetry. In SOI, this can either be done by converting the layer structure to a membrane with a selective oxide etch, or by depositing oxide on top of the Silicon core. Alternatively, one can etch the photonic crystal holes deeply into the oxide, and thus replace the homogeneous oxide cladding by a cladding with a large fraction of air.

4.6.3 Band Shifting

Another loss mechanism can occur because of inhomogeneities in the fabricated structures. In a perfect photonic crystal waveguide made in a layer structure with high index contrast, the wavelength range of the guided mode below the light line is fixed in place. As we have seen in the previous chapter, this wavelength range is often very narrow in the case of SOI.

However, in a real structure, all holes of the photonic crystal structure won't be identical. Also, depending on the fabrication technique, their position will not always be exactly on a lattice point. As a re-

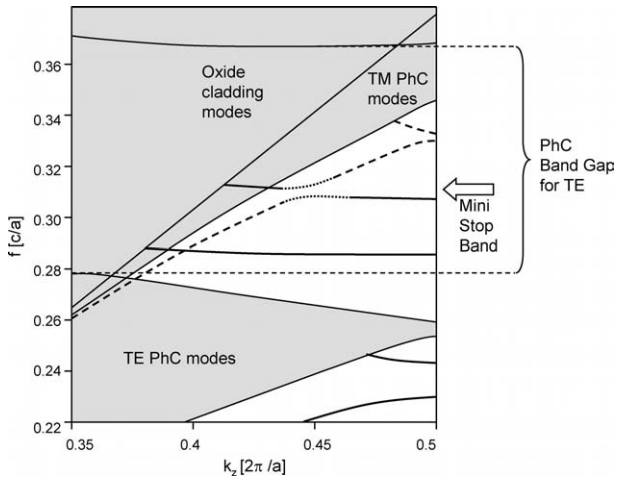

Figure 4.27: TE/TM interaction in SOI, with the creation of a mini-stop-band. In this detail of the full band diagram, the quasi-TE modes are printed in full line. The index-guided TM mode is printed in dashed line. Where the modes interact, the resulting mixed modes are neither TE or TM.

sult, the transmission band of the waveguide mode can shift or distort slightly along the waveguide. When the wavelength range where the mode is located below the light line is very narrow, this effect can locally couple the guided mode to the radiation modes. As this coupling is very strong when using a high vertical index contrast, this can introduce significant losses.

To alleviate this loss mechanism, both wide-band, tolerant designs and reproducible, high-precision fabrication tools are needed.

4.7 Conclusion

In this chapter, we have studied some loss mechanisms in nanophotonic waveguides in general, and photonic crystal slabs in particular. When the photonic crystal waveguide mode is above the light line of the slab cladding, the light can couple to the radiation modes. This coupling increases for higher refractive index contrast. When operating below the light line, this coupling does not occur as long as the periodicity of the waveguide is not broken. Basically, there are two favourable regimes to reduce these intrinsic scattering losses. One can either use very low refractive index contrast, where both waveguides and defects radiate a small amount of the propagating light. This regime can be

useful when one wants to make small components with a large number of defects. Alternatively, one can use a very high vertical index contrast and design a waveguide with a guided mode below the light line. This waveguide is intrinsically lossless, but can scatter strongly at defects like bends and splitters.

Both photonic crystal slab waveguides and photonic wires are sensitive to additional losses due to substrate leakage and scattering at sidewall roughness. Substrate leakage is easily avoided by choosing an adequately thick buffer layer between the waveguide core and the substrate. For scattering at sidewall roughness, there is little that can be done except for process optimalisation in order to avoid irregularities in the first place. However, we have calculated that for photonic crystal slabs, scattering at roughness is much lower when using a high vertical index contrast instead of a low vertical index contrast.

With this knowledge of loss mechanisms in both photonic wires and photonic crystal slab waveguides we can deduce a number of conditions for good waveguide designs.

When choosing a low refractive index contrast, waveguide operation is mostly restricted above the light line because of the high refractive index of the cladding: the light line is not very steep and the PBG of the lattice is located in the radiation continuum, as can be seen in figure 2.24 on page 41. For conventional III-V heterostructures, it is therefore beneficial to keep the vertical index contrast as low as possible to minimize out-of-plane scattering. However, a low contrast implies a widely spread guided mode. In order to avoid additional scattering at the bottom of the etched holes, deep etching technology is required. Also, the buffer layer between the core and the substrate should be thick enough to avoid substrate leakage. However, one major drawback of low vertical index contrast is the larger scattering losses due to roughness. One can only alleviate this with sufficiently optimised processing.

When working with high vertical index contrast, and therefore a low cladding index, the light line of the cladding is much steeper and the photonic bandgap of the lattice shifts (partially) below the radiation continuum. It is now possible to design a structure which has a bandgap-guided mode below the light line. Such a mode is intrinsically lossless in a straight waveguide. However, defects should be carefully engineered using 3-D modelling tools in order to minimise the scattering. One should also take care that the waveguide mode is guided (both in the PBG and below the light line) in a sufficiently broad wavelength range.

Waveguide materials with high vertical index contrast, like SOI and membranes, are less prone to scattering at sidewall roughness. However, optimised processing is still a requirement for low-loss waveguides.

For both photonic wires and photonic crystal slabs we have found that the TE polarisation is less sensitive to sidewall roughness than the TM polarisation. This is an advantage, as the TE polarisation has a photonic bandgap for a lattice of air holes etched in a high-index material. This configuration is much easier to fabricate than a structure consisting of high-index pillars, which has a PBG for the TM polarisation.

Chapter 5

Fabrication with CMOS Technology

"So you really need high-resolution lithography?"

Chapter 5

Fabrication with CMOS Technology

As already shown in chapter 3, nanophotonic structures need to be fabricated with very high accuracy and reproducibility. Additionally, the tolerances on sidewall roughness are very strict, as discussed in chapter 4. This chapter will explore the use of advanced CMOS technology, including deep UV lithography, for the fabrication of nanophotonic components.

5.1 Nanophotonic Fabrication Technology

Nanophotonic structures have a typical smallest feature size of a few hundred nanometers. However, the accuracy with which the structures have to be fabricated is of the order of $10nm$ or even less, especially when wavelength-selective functions are involved.

Many conventional photonic components are often fabricated with a combination of lithography and etching: the patterns of the structures are defined into a sensitive resist layer. After development, the pattern is transferred from the resist into the device material by an etching process.

5.1.1 Pattern Definition

The definition of the patterns into the resist can be done in several ways. When the resist is photosensitive (so-called photoresists), a pre-defined mask can be imaged onto the photoresist. This can either be done using

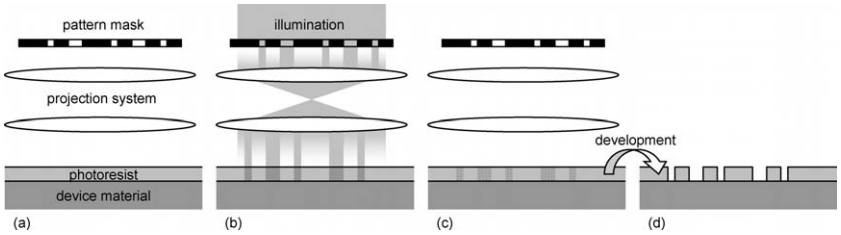

Figure 5.1: Optical projection lithography. (a) setup, (b) imaging of the pattern, (c) the pattern is latent in the photoresist, (d) after development.

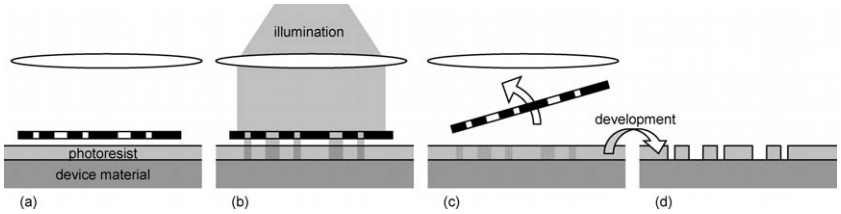

Figure 5.2: Optical contact lithography. (a) setup, (b) imaging of the pattern, (c) the pattern is latent in the photoresist, (d) after development.

a projection system (*projection lithography*: figure 5.1), or by positioning the mask right on top of the photoresist (*contact lithography*: figure 5.2). The resolution of this optical lithography is dependent on a number of factors. The main limitation is imposed by the illumination wavelength: The shorter the wavelength, the smaller the features that can be defined. For conventional photonic structures, wavelengths in the optical or near UV range are used, the shortest wavelength for projection lithography being $365nm$. Typically, this limits the feature size for dense features to approximately $600nm$ or more, depending on the optics. For nanophotonic structures, this resolution is insufficient.

Alternatively, the patterns can be written directly into the resist. With *e-beam lithography* (figure 5.3), a focused electron beam scans over the resist. Because high-energy electrons have a much shorter wavelength it is possible to define much smaller features. However, all features have to be written serially, and this makes e-beam lithography a slow process. On the other hand, because no predefined mask is needed, the technique is very flexible, especially for research purposes.

An alternative process that is steadily gaining ground in nanophotonics is the use of *nano-imprint lithography* (NIL) (figure 5.4). Here, a predefined mold is used to stamp the pattern into a resist layer [91]. The

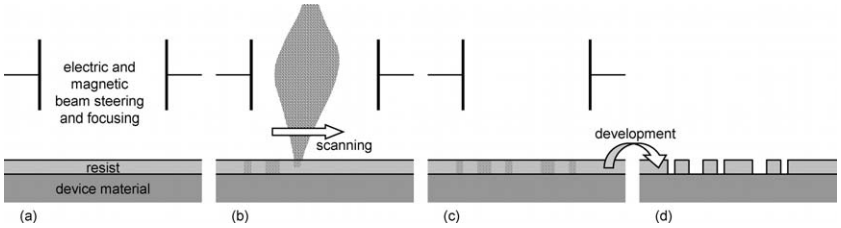

Figure 5.3: E-beam lithography. (a) setup, (b) writing of the pattern with a scanning electron beam, (c) the pattern is latent in the resist, (d) after development.

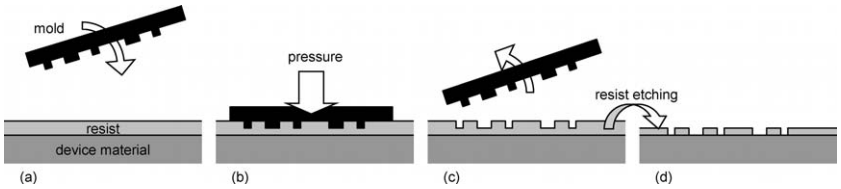

Figure 5.4: Nano-imprint lithography. (a) A predefined mold with the negative image if (b) pressed into the resist. (c) When the mold is removed, a partial imprint remains in the resist. (d) After a partial etch, the resist can be used as a mask.

predefined mold can be fabricated once with a slow, high-resolution process like e-beam lithography.

Deep UV lithography differs from conventional optical lithography in the use of the wavelength. With wavelengths of $248nm$, $193nm$ and in the future $157nm$[1], much higher resolutions can be reached. However, this comes at a considerable cost: The high-quality projection optics for deep UV wavelengths are extremely complex. The price tag of a deep UV stepper calls for a large initial investment, making the technology only available to industries with a sufficient market potential, as is the case with advanced CMOS components. For these large-volume applications, deep UV lithography is the most attractive technology because of its high throughput and acceptable working costs.

We explored the possibilities of using deep UV lithography for the fabrication of photonic components. If this technology is to become a viable alternative to e-beam lithography or NIL, a large-volume market

[1]Even further in the future, extreme UV lithography will use wavelengths in the order of $15nm$. At these wavelengths, reflective optics, including reflective masks, are required. [92]

for nanophotonic components is required to justify the costs of setting up a deep-UV-lithography-based manufacturing process.

5.1.2 CMOS Processes

To satisfy the demand for more computer power, the CMOS industry continues to shrink the basic building blocks of integrated circuits. This increases switching speed, reduces the required operating voltage and allows the integration of more components on a chip. Transistor gate widths, typically one of the smallest features on a CMOS chip, in commodity chips have fallen from $350nm$ down to $130nm$ at the time of writing to $90nm$ in the near future. Research to shrink this dimension down to $65nm$ and $40nm$ is already ongoing [93].

To fabricate these small structures reliably in sufficient quantities, deep UV lithography is used to define the patterns. The cost of the lithography rises significantly for shorter illumination wavelengths[2], and therefore as many features as possible are defined at $248nm$ wavelength. Only for the smallest features, illumination at $193nm$ is used. Projection lithography at $157nm$ is under development [94].

Also, as CMOS structures are made on $200mm$ or $300mm$ wafers, sufficient reproducibility and homogeneity of all processing steps is needed to guarantee an acceptable yield. The wafers are not illuminated in a single step, but in a stepper which prints the same die in a grid-like pattern on the wafer. Typically, a die can be as large as $25 \times 30mm^2$. The mask contains the layout of a single die, but at a scale 4 times larger than on the wafer. The $4X$ reduction is done by the optical projection system.

In this work, we have explored the use of deep UV lithography to make nanophotonic structures in Silicon-on-insulator. For the fabrication, a cooperation was set up with IMEC in Leuven to make use of the CMOS fabrication processes. We made use of deep UV lithography at $248nm$ and $193nm$, dry etching of both silicon and oxide, and in a later stage, plasma treatment of the resist. For process evaluation, top-down SEM inspection is possible on-site.

[2]This is because the range of materials to develop high-quality optics and photoresists gets smaller and smaller for shorter wavelengths

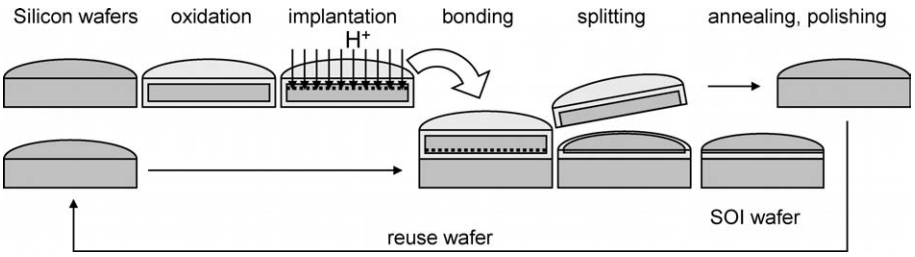

Figure 5.5: UNIBOND® fabrication of SOI wafers.

5.1.3 Silicon-on-Insulator

As we have already mentioned in the previous chapters, SOI is a very good material system for nanophotonic components [53]. SOI waveguides with thick SOI cores have already been around for more than a decade. However, while these large-core waveguides can interface easily with optical fibres, they are not suited for nanophotonics, as the lateral confinement is not very strong [95]. Also, in many cases these waveguides are multimode [96]. Apart from being an good material for photonic waveguides, SOI is also compatible with CMOS processes, and commercially available in $200mm$ wafers.

High-quality Silicon-on-insulator wafers are typically fabricated using wafer bonding. For our experiments, we used commercial wafers from *SOITEC* fabricated with the *UNIBOND*® process (figure 5.5).

First a wafer is oxidised to create the buried Oxide layer. Then Hydrogen ions are implanted at a well-controlled depth, creating a Smart Cut® [97]. This wafer is bonded to a clean Silicon wafer. The substrate of the first wafer can now be separated along the implanted Smart Cut®, and then annealed and polished.

For the first experiments, we used standard UNIBOND® wafers with a buried oxide of $400nm$ and a top silicon layer of $205nm$. Because this oxide thickness is too thin, as discussed in chapter 4, causing optical leakage to the substrate, we switched to custom-made wafers with a Silicon thickness of $220nm$ and an oxide layer of $1\mu m$. As seen in chapter 4, this buffer thickness provides adequate isolation from the substrate for the TE polarisation. The thickness of the oxide is limited to $3\mu m$ by the Smart Cut process employed by SOITEC.

Note that SOI is not only a good material for photonic applications. For high-end electronics, SOI is used to better isolate the active structures from the Silicon substrate, eliminating carrier leakage and reduc-

ing capacitive coupling. However, for these applications, both the top Silicon layer and the buried oxide are much thinner than for nanophotonic waveguides.

5.1.4 The Fabrication Process

The fabrication process with deep UV lithography is similar to that of conventional optical projection lithography. The process flow is illustrated in figure 5.6 [49]. First, the photoresist is coated on top of a $200mm$ SOI wafer, and then pre-baked. On top of the resist, an anti-reflective coating is spun to eliminate reflections at the interface between the air and the photoresist. These reflections increases standing waves in the photoresist, and therefore inhomogeneous illumination. Then the wafer is sent to the stepper, which illuminates the photoresist with the pattern on the mask. As a 200mm wafer can contain many structures, the die with the pattern is repeated across the wafer. This can be done by varying the exposure conditions (focus and exposure dose), which makes it possible to do detailed process characterisation. After lithography, the resist goes through a post-exposure bake, and is then developed. For our experiments we used Shipley UV3 resist. The developed photoresist is then used directly as a mask for etching. We can either etch just the top Silicon layer, or etch through both the Silicon and the oxide layer. Also, a number of post-processing steps are possible, including thermal oxidation or oxide deposition.

5.1.5 Nanophotonics versus CMOS

While deep UV lithography is capable of printing features with the dimensions of photonic wires and photonic crystals, there are some significant differences between nanophotonic components and typical CMOS components: CMOS components are layered structures. Each layer contains only critical structures of a certain type (transistor gates, contact holes, ...), so the process can be optimised for each layer individually. In planar nanophotonics, all structures are fabricated on the same level, and small alignment tolerances require that all structures be fabricated in the same lithography step. Photonic crystals and photonic wires are very different structures, and the optimal process conditions differ between them. Therefore, process optimisation for nanophotonic components will have to make different compromises than for CMOS structures. As we will see further, it is not straightforward to fabricate both photonic wires and photonic crystals on target at the same time

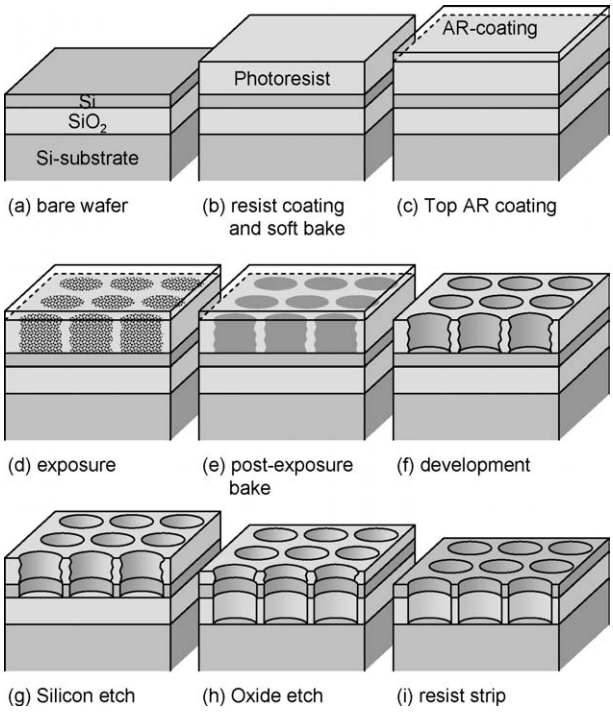

Figure 5.6: The process flow for the fabrication of nanophotonic structures in SOI.

Also, the requirements for sidewall roughness are very different in nanophotonics and in CMOS. In the former, all sidewalls should be kept smooth to reduce scattering, while in CMOS structures, only in narrow lines the effect of line-edge roughness is felt in the series resistance of the line.

5.2 Deep UV Lithography

Because lithography at a wavelength of $248nm$ is now the mainstream fabrication tool for high-end CMOS, we explored its possibilities for the fabrication of nanophotonics. For this, we used an *ASML PAS5500/750* stepper/scanner connected to an automated track for preprocessing (coating and baking) and postprocessing (baking and developing).

5.2.1 Resolution in Optical Lithography

As mentioned before, the resolution of optical lithography is dependent on a number of factors [50]. For projection lithography, the wavelength and the properties of the optical imaging system are the most important. Figure 5.7 shows a schematic of the optical system in a deep UV stepper.

The *numerical aperture* (NA) of the projection system is defined by the sine of the largest angle under which a beam can still pass through the optical system. In figure 5.8 the middle of the 3 beams is just not blocked by the circular aperture. The NA of this lens system is therefore

$$NA = \sin\left(\theta_{NA}\right), \tag{5.1}$$

with θ the angle of the beam with the optical axis in the image plane. The position of the circular aperture that limits the NA of the lens system is the NA plane. For a uniform illumination of the wafer, the projection system is conceived in such a way that the *NA plane coincides with the image plane of the light source*. This is called a *Köhler* illumination setup.

Perfectly Coherent Light

To get an idea of the resolution of an optical projection lithography system, let's consider the light source to be a point emitting perfectly coherent light. This, with suitable collimator optics, creates a perfectly

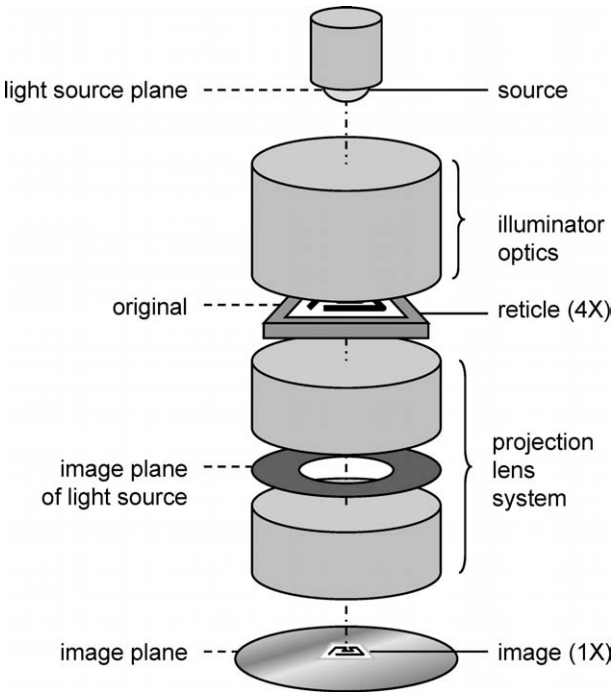

Figure 5.7: Schematic of the optics in a deep UV stepper use a Köhler illumination system. The image of the light source is in the same plane as the numerical aperture. This results in a uniform wafer illumination. Also, the projection optics, consisting of at least 2 lens systems, have a reduction factor of $4X$.

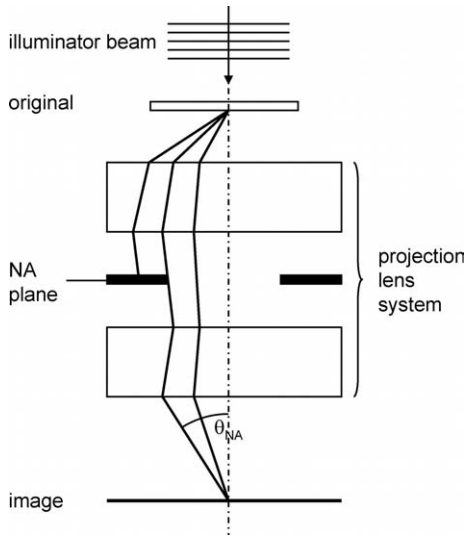

Figure 5.8: The numerical aperture in the projection system is given by $\sin(\theta_{NA})$. θ_{NA} is the angle of the beam with the optical axis in the image plane for the beam that can marginally pass through the lens system. The plane in the lens system that is the limiting factor for this angle is the NA plane.

homogeneous plane wave incident on the mask. If such a beam illuminates an isolated structure, like a pinhole or a line, as shown in figure 5.9a, the light is diffracted omnidirectionally, as dictated by the Huygens principle. When this light is passed through a lens system, the light scattered at high angles with respect to the axis will be blocked by the aperture in the NA plane of the projection lens system. The narrower the slit, the more light is diffracted at higher angles, and the less light is available to reconstruct the image of the slit in the image plane. Therefore, the image of the slit will gradually degrade as the slit width w gets smaller with respect to the illumination wavelength λ. Typically, the image can be reconstructed when

$$w > \frac{\lambda}{2NA}. \tag{5.2}$$

This is not a hard limit, and the resolution of isolated features is very much dependent on the resist and the processing. While degradation of the image of an isolated feature is continuous, this is not true for densely packed, and especially periodic, structures. In figure 5.9b we see that each slit of a grating with period a radiates omnidirectional. However, these individual contributions will interfere destructively ex-

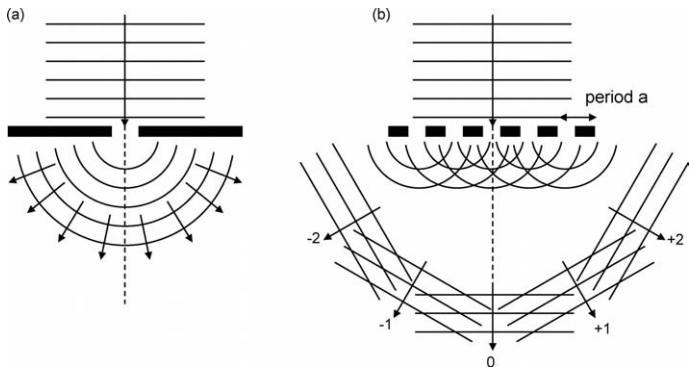

Figure 5.9: Diffraction of an isolated slit and a periodic array of slits. (a) At an isolated slit, the light is diffracted omnidirectional, according to the Huygens principle. (b) At a periodic grating of slits, interference directs the diffraction under specific angles.

cept under specific angles. So instead of a beam with a continuous angular distribution, we end up with a collection of diffraction order consisting of plane waves at discrete angles θ_i:

$$\sin(\theta_i) = i\frac{\lambda}{a}, \tag{5.3}$$

with i an integer number.

When this diffraction pattern is passed through the projection lens system, only a discrete number of diffraction orders with $\theta_i < \theta_{NA}$ can pass through the aperture in the NA plane. To reconstruct the original image, at least a few diffraction orders are needed. This is illustrated in figure 5.10, where we plotted the imaging of a periodic structure through a projection lens system. The plane wave coming from the point light source is diffracted by the periodic structure and split up into a plane wave for each of the discrete diffraction orders. In the NA plane, which coincides with the image plane of the light source, these plane waves will each form an image of the light source, i.e a point. The larger the diffraction angle of the plane wave, the further away from the optical axis the source image will be formed. For the top structure, with the large period a_1, the image of the light source in diffraction orders 0, ± 1 and ± 3 will fit within the aperture in the NA plane. Therefore, the image of the periodic structure will be reconstructed from these diffraction orders. For the smaller period a_2 in figure 5.10b, only diffraction orders 0 and ± 1 can pass through the NA plane. The resulting pattern is therefore only sinusoidal, and not so well defined as with the larger

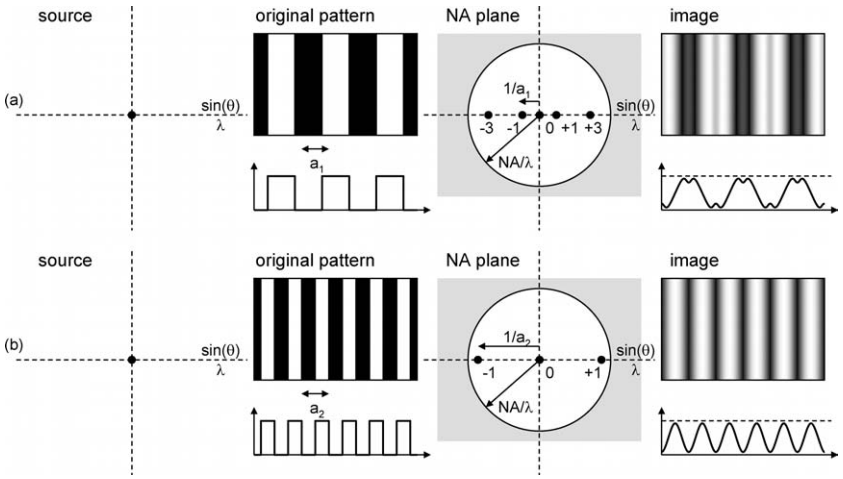

Figure 5.10: Imaging of a periodic structure with a perfectly coherent beam. The patterns in the various planes of the projection system are shown. Left to right: The light source, the original image (mask) with intensity profile, the NA plane with the aperture, and the resulting image on the wafer. The situation is plotted for a structure with (a) a large pitch, where the orders 0, ± 1 and ± 3 can pass through the NA, and (b) for a smaller pitch, where only the diffraction orders 0 and ± 1 can pass through the lens system.

period. Note that each position in the plane of the light source and the NA plane corresponds with a plane wave at an angle θ in the mask plane and image plane on the wafer and vice versa[3].

When only the zeroth order plane wave is transmitted, the result is a homogeneous intensity in the image plane. When both order -1 and +1 can pass, a sinusoidal intensity modulation is added. For higher diffraction orders, additional intensity modulation at higher spatial frequencies is added. Therefore, in order to print a periodic structure, at least the diffraction orders 0 and ± 1 should pass through the lens. This is the case when

$$a > \frac{\lambda}{NA}. \tag{5.4}$$

Unlike with isolated structures, this is a hard limit, below which no imaging is possible, irrespective of the resist and the processing. However, as we will see next, using a different light source can increase the resolution.

[3]As we use optics with a 4X reduction factor, the angles in the image plane are different that those in the mask plane. If the reduction factor would be 1X, the angles would be the same. In this text, we use the angle θ in the image plane.

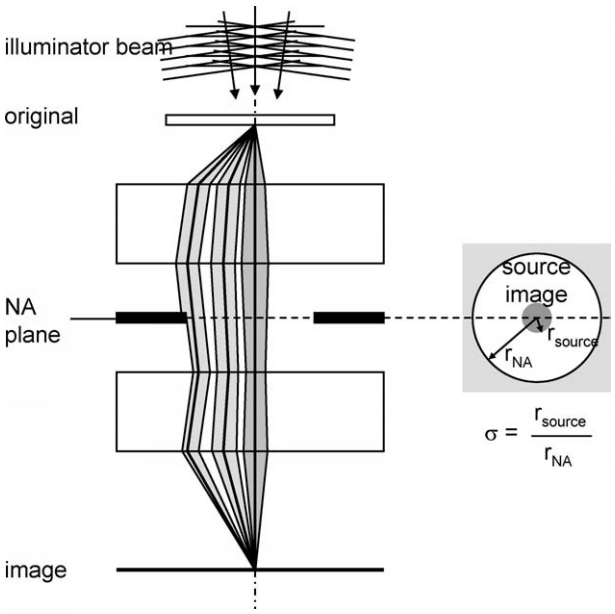

Figure 5.11: Projection system with partially coherent light. Instead of a illumination with a plane wave, the source has a finite size, resulting in an angular distribution in the plane of the mask. The coherence of the light is described by the fill factor σ, which is the ratio of the radius of the source image in the NA plane and the numerical aperture.

Partially Coherent Light

When we replace the point light source by a light source with a finite size, the mask will no longer be illuminated by a single plane wave propagating along the optical axis, but by a beam with a certain angular distribution. This incident light will be diffracted by the patterns on the mask and then passed through the projection system, also with a certain angular distribution. The source image in the NA plane is now also finite in size. The ratio between the source image radius and the numerical aperture is called the *spatial coherence factor* σ of the illumination [98, 50]:

$$\sigma = \frac{r_{source}}{r_{NA}}. \tag{5.5}$$

This is illustrated in figure 5.11.

For periodic structures, each diffraction order now consists of a beam with an angular distribution, instead of a single plane wave at a discrete angle. Therefore, in the NA plane, the different diffraction

orders will no longer be points, but also images of the source. In figure 5.12 we have again plotted the transmission of a periodic grating, but now with a partially coherent light source. In figure 5.12a, we have taken the period a_2 from 5.10b. We see now that the ± 1 diffraction orders are partially blocked in the NA plane. This results in a reduced image contrast. On the other hand, in figure 5.12b, we have further reduced the period. Now, only a small fraction of the ± 1 diffraction orders is passed through the aperture, and the image contrast is even further reduced. However, imaging is still possible, while with perfectly coherent light, the ± 1 diffraction orders would have been completely blocked. This shows that partially coherent light can be used to attain resolutions beyond those possible with perfectly coherent light, but at the cost of lower image contrast. For a beam with wavelength λ and spatial coherence factor σ, the minimum period a where part of the ± 1 diffraction order is still transmitted is

$$a > \frac{\lambda}{NA\,(1+\sigma)}. \tag{5.6}$$

Practical Resolution Limit

For both a perfectly coherent and a partially coherent source, the resolution is determined by the period of the grating, and not by the width of the individual lines. The resolution can be improved by going to shorter wavelength, increasing the numerical aperture and by increasing the spatial coherence factor σ. However, these techniques have drawbacks. Shorter wavelengths (from $248nm$ to $193nm$ to $157nm$) require highly advanced optics in materials suitable for deep UV wavelengths. Similarly, optics with sufficient imaging field, high numerical aperture and little aberrations are very expensive. Also, working with a large NA implies a very limited depth of focus in the imaging plane, the light comes in at steep angles. Increasing the size of the light source is easier, but while this increases the resolution for small periodic structures, it also decreases the image contrast of all other structures.

While the wavelength, numerical aperture and the shape of the source dictate the maximum theoretical resolution, the real resolution is very dependent on the resist and the processing. In practice, the minimum period a calculated in formulas 5.4 and 5.6 should be multiplied by a factor $k > 1$. The better the resist and processing, the closer k approaches 1.

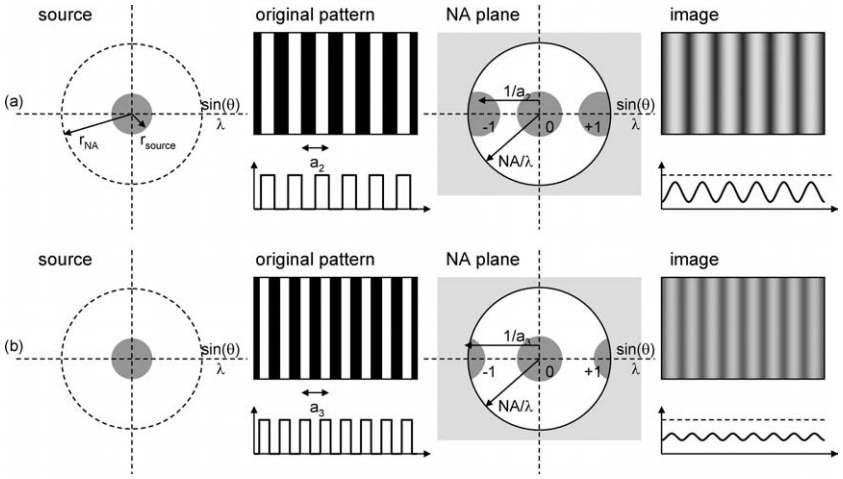

Figure 5.12: Imaging of a periodic structure with a partially coherent beam. The patterns in the various planes of the projection system are shown. Left to right: The light source, which now has a finite size, the original image (mask) with intensity profile, the NA plane with the aperture, and the resulting image on the wafer. The situation is plotted for a structure with (a) a relatively pitch (same as in figure 5.10b, where the orders 0, ± 1 can pass through the NA-plane, and (b) for a smaller pitch, where only the diffraction orders ± 1 can are only partially transmitted. However, with perfectly coherent light, the orders ± 1 would not be transmitted at all.

A layer of photosensitive resist has a nonlinear response to the illumination intensity. In an ideal case, the material response is a step-function. For older technology, where imaging near the resolution limit was not an issue, photoresists had a less nonlinear response. Current chemically amplified photoresists perform much better in that respect.

When only a limited number of diffraction orders are imaged, the intensity in the resist is a sinusoidal curve. In combination with a good resist, the areas where the intensity is higher than the resist threshold will react and the other areas won't. As the peak intensity depends on the power of the light source, the exact transition between 'light' and 'dark' areas is determined by the exposure dose. In the case of our periodic grating, increasing the exposure dose will therefore change the duty-cycle of the grating, reducing the dark areas and increasing the light areas. The energy range in which the structures are considered to be printed on-target is called the *exposure latitude* (EL). The more nonlinear the response of the resist, the more accurate the features can be defined.

A layer of photoresist has a finite thickness, sometimes more than $500nm$. Therefore, the projected image should be homogeneous in a layer that is several wavelengths thick. When operating with high numerical aperture, the *depth of focus* (DOF) is very small, and the imaging will not be homogeneous throughout the resist.

Together, the exposure latitude and the depth of focus define the *process window*. For most CMOS structures , the process window is the parameter range in which the printed structure varies within 10% of the target size. However, in nanophotonic applications, the maximum allowable deviation is of the order of 5%. For reproducible fabrication, the process windows should be sufficiently large. As we will see further, the process window decreases dramatically as the features shrink closer to the theoretical resolution.

An additional problem associated with the finite thickness of the resist is standing waves. The incident light is reflected both at the top and the bottom interface of the resist. Because of these reflections, a cavity is formed, creating standing waves in the resist. The resulting intensity modulation gives rise to uneven exposure throughout the resist. There are two solutions to this problem, which, if necessary, can be combined. To remove the reflections at the top interface, a thin anti-reflective coating is deposited on top of the resist, a so-called *top anti-reflective coating* (TAR). This coating has a refractive index that is the geometric average

of the refractive index of air and the photoresist. The thickness of the layer is a quarter of the illumination wavelength in the material.

The alternative solution is to absorb the light at the bottom of the resist layer, so that it isn't reflected. However, while very effective, this *bottom anti-reflective coating* (BARC) is not developed together with the resist. Therefore, before etching the underlying structures, the BARC has to be etched first. Because this etching consumes a part of the resist, the original resist layer has to be thicker than without BARC, and this in turn reduces the depth of focus. In structures where the wafer is not flat, e.g. in the case where there are already structures defined on it, the BARC is more effective than the TAR. Because for planar nanophotonics there is little topography on the wafer, we found a TAR to suffice.

For the majority of our experiments, we used an illumination wavelength of $248nm$, a numerical aperture $NA = 0.63$ and a spatial coherence factor $\sigma = 0.4$. With these settings, the theoretical minimum period that can still be imaged is

$$\begin{aligned} a_{min} &= \frac{\lambda}{NA\,(1+\sigma)} \\ &= \frac{248nm}{0.63\,(1+0.4)} \\ &= 281nm. \end{aligned} \tag{5.7}$$

However, in practice, the minimum period will be somewhat larger. As we will see, we have experimentally fabricated periodic structures with a pitch of $400nm$ with a sufficiently large process window.

5.2.2 Photonic Crystals versus Contact Holes

Even though the resolution of deep UV lithography and associated CMOS processes is adequate for most nanophotonic structures, there are some considerable differences between CMOS structures and nanophotonic waveguides.

As we have seen in chapter 2, the optimal photonic crystal lattice for TE polarisation is a triangular lattice of air holes where the holes have a large fill factor, i.e. the hole diameter is a significant fraction of the pitch. The best equivalent in CMOS is a dense array of contact holes used to connect the different metal interconnect layers. These contact holes are typically arranged in a square lattice, and the hole diameter never exceeds half of the pitch. This is illustrated in figure 5.13.

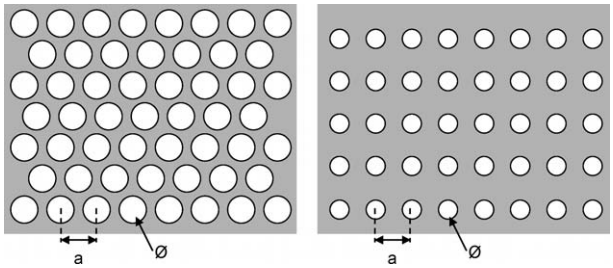

Figure 5.13: Photonic crystal versus contact holes. Left: Photonic crystal, with a triangular lattice and a $\oslash/a > 0.5$. Right: dense contact holes, with a square lattice and a $\oslash/a \leq 0.5$.

5.2.3 First Experiments: Dense Square Lattices

As photonic crystal structures are unlike CMOS structures, there were initially no masks with suitable test patterns available. However, we could make use of the CONCLEAVE2000 mask, which contains test patterns for contact holes arranged in dense square lattices. The hole shape on the mask is square, but it will be imaged as a circular hole, as higher-order diffraction orders cannot pass through the lens. As already explained in the previous paragraph, these structures resemble photonic crystals, but with some important differences. One of the important aspects of photonic crystal fabrication that we wanted to look into was the possibilities of making superdense lattice, i.e. where the hole diameter is larger than the spacing in between. In the test patterns on CONCLEAVE2000, this was not the case.

However, by increasing the exposure dose, we can increase the light areas on the wafer, and therefore create a superdense lattice. Figure 5.14 shows some resist patterns of a square lattice of holes with a pitch of $500nm$ and a hole diameter (on the mask) of $250nm$. Using overexposure, the hole diameter was increased to $330nm$. In the cross section, we can see that the resist pattern has a very vertical profile.

Figure 5.15 shows the printed hole size as a function of exposure dose. As we can see, overexposure makes it possible to make the holes substantially larger than their original design size. Also, the exposure latitude is more than adequate for most structures.

5.2.4 Photonic Crystals: Superdense Triangular Lattices

While the first tests with the CONCLEAVE2000 mask showed that it is indeed possible to fabricate superdense lattices, more representative

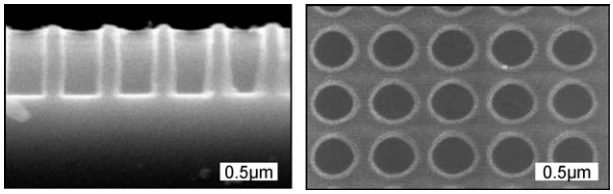

Figure 5.14: Resist patterns of a superdense square lattice of holes. Left: Cross-section. Right: Top-down. The lattice pitch is $500nm$, the hole diameter is $330nm$. The resist thickness is $800nm$.

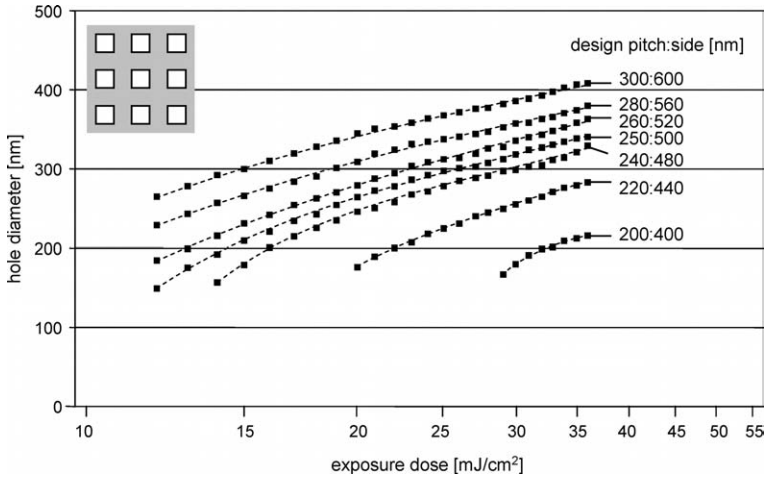

Figure 5.15: Size of a square lattice of holes as a function of exposure dose for different design pitches and diameters. The design side:pitch ratio is always 0.5.

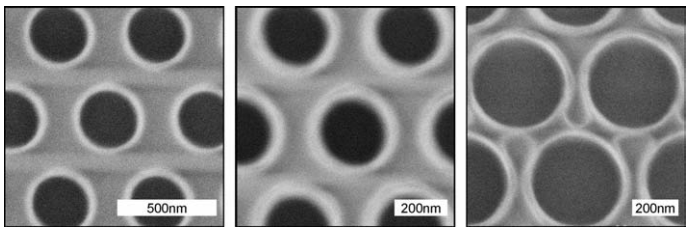

Figure 5.16: Resist patterns of a superdense triangular lattice of holes. Left to right: $300nm$ holes with $500nm$ pitch, $240nm$ holes with $400nm$ pitch, $420nm$ holes with $500nm$ pitch. In the last case, the resist between the holes is already partially illuminated.

test structures were needed. Typical photonic crystals consist of superdense triangular lattices. To this purpose we designed a test mask, *PICCO_1*, which contains a variety of nanophotonic test structures, including triangular lattices with varying pitch and fill factor. Instead of square holes, we now used hexagonal holes on the mask, in order not to break the symmetry of the lattice.

Figure 5.16 shows some SEM pictures of resist patterns of such superdense lattices. For relatively large pitches around $500nm$, the imaging is very good. For a pitch of $400nm$, near the resolution limit of the stepper, the quality is somewhat less, but still acceptable. In the situation where the hole is very large compared to the pitch, the resist in between the holes is already partially illuminated, as can be seen in the right part of figure 5.16.

In these extremely dense structures, the hole size is no longer the limiting factor, but the linewidth between the holes. As this line is very narrow, the exposure latitude of these structures is drastically reduced. This can be seen in figure 5.17, which plots the printed hole size as a function of exposure dose for varying pitches and design diameters. The holes with a designed diameter:pitch ratio of 0.6 clearly allow the the largest range of printed hole diameters while still keeping a sufficiently large exposure latitude, even for the smallest pitch of $400nm$.

5.2.5 Narrow lines: Isolated Lines

Photonic wires are a very different type of structures than photonic crystals, and have different optimal process conditions. For CMOS devices, narrow lines and dense holes are not printed together. But for nanophotonic structures, this is necessary to avoid alignment prob-

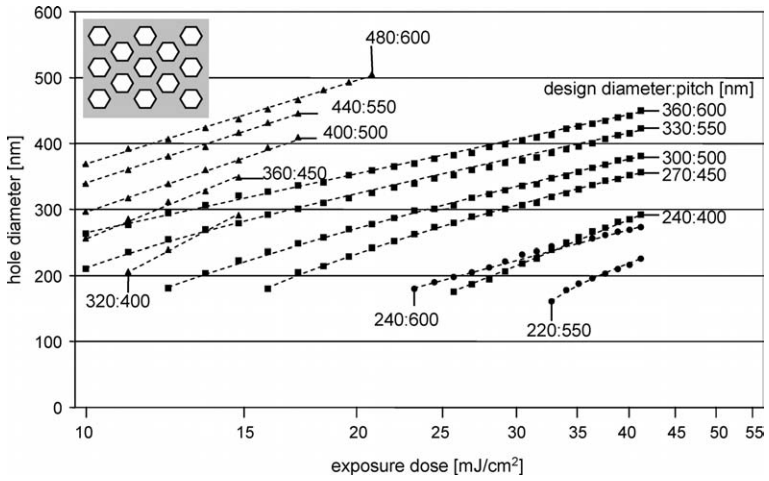

Figure 5.17: Size of a triangular lattice of holes as a function of exposure dose for different design pitches and diameters. The designed diameter:pitch ratio is 0.4 (circles), 0.6 (squares) and 0.8 (triangles).

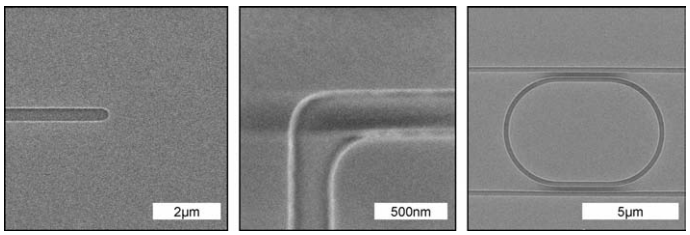

Figure 5.18: Resist patterns of isolated lines. Left to right: Line end, right angle, racetrack resonator.

lems. We therefore conducted experiments to print narrow isolated lines with the process we used for the dense lattices of holes.

Figure 5.18 shows some resist patterns of photonic wires printed with the same process as the photonic crystals. Even though this process is not optimised for such structures, both the quality and the process window is sufficient for the typical line widths of photonic wires ($400nm$ - $600nm$). This is also illustrated in figure 5.19, which plots the line width as a function of exposure dose for various design widths.

5.2.6 Combining Lines and Holes

As already mentioned, one of the difficulties of fabricating nanophotonic components is the requirement to print both photonic wires (iso-

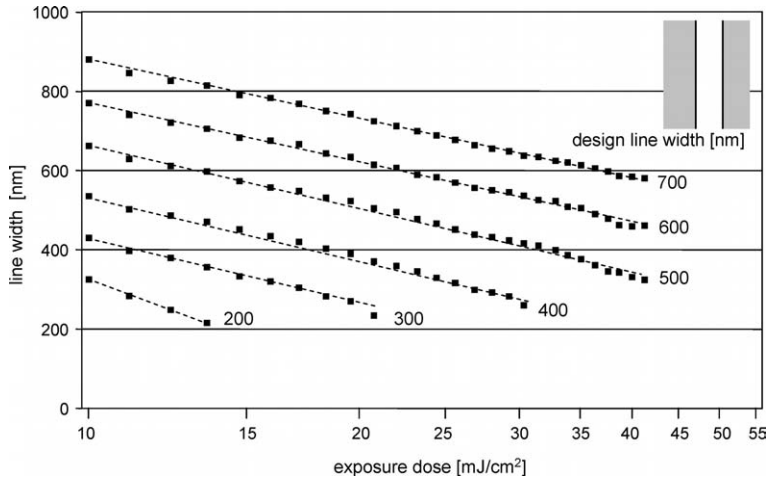

Figure 5.19: Width of isolated lines as a function of exposure dose for different design widths.

lated lines) and photonic crystals (superdense lattices of holes) together in the same lithography step. As we can deduce from the graphs in figures 5.17 and 5.19, the dose-to-target for lines and holes is quite different. Figure 5.20 shows the bias for an isolated line printed at the dose-to-target of a superdense lattice of holes with $500nm$ pitch and $300nm$ hole diameter. As this dose-to-target is higher than the dose-to-target for the isolated line, the lines will be overexposed and therefore printed narrower than originally designed.

In order to print the lines correctly, a bias needs to be applied in the design to either the holes or the lines to print both together on target. Because it is easier to change the design size of an isolated structure, the bias is best applied to the lines. Figure 5.21 shows the design line width needed to print a specific target line width at a given dose. As we can see, at the dose of $25mJ/cm^2$ where $300nm$ holes with $500nm$ pitch print correctly, a $50nm$ bias needs to be applied to a $500nm$ line to print it on target. This correction should be known in advance, because it needs to be applied directly on the mask. As masks for deep UV lithography are very expensive, a good process characterisation is needed.

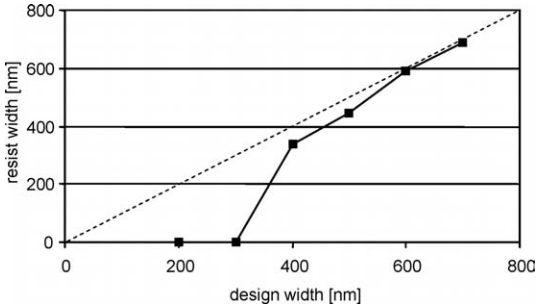

Figure 5.20: Measurements of Line-hole bias. Resist line width as a function of design line width of a narrow line when printed at the dose-to-target of a superdense triangular lattice ($500nm$ pitch, $300nm$ hole size). The dashed line indicates features printed on target.

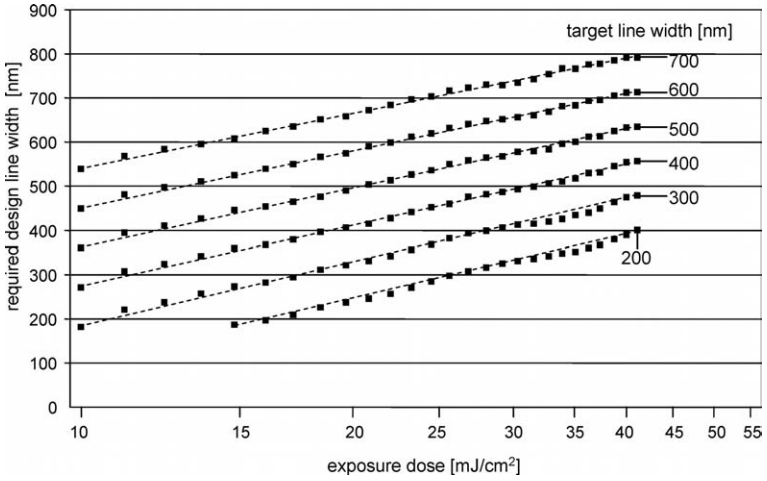

Figure 5.21: Required design line width for specific target line width as a function of exposure dose.

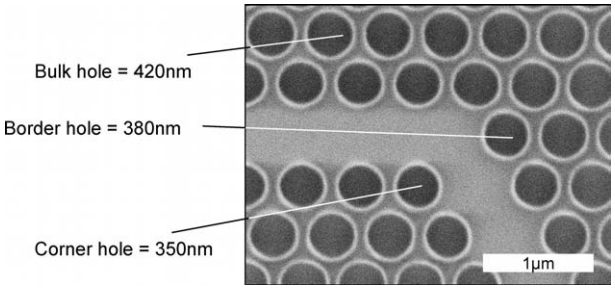

Figure 5.22: Example of optical proximity effects. The holes near the photonic crystal defect are printed smaller than the holes in the bulk of the lattice. The lattice pitch is $530nm$.

5.2.7 Optical Proximity Effects

Photonic crystals are superdense periodic structures with feature sizes close to the illumination wavelength. This causes the images of neighbouring holes to interfere during lithography. Because of this, holes in a photonic crystal may interfere constructively and print larger or interfere destructively and print smaller than semi-isolated holes. In uniform lattices, this effect is not noticeable, as the illumination energy will be chosen to print the holes in the lattice on target. However, at the boundaries of the lattice, or near defects like a waveguide or cavity, some holes lack neighbors and will therefore print differently than their counterparts in the bulk of the lattice. This phenomenon is called *optical proximity effects* (OPE). An example is given in figure 5.22: the holes near the line defect are $40nm$ smaller than the holes in the bulk, and in the corner this effect is even worse, with the corner hole being $70nm$ smaller.

The functionality of photonic crystal waveguides and related components is largely determined by the holes near the (line) defects. Figure 5.23 illustrates the effect of OPE on the guided mode of a simple W1 photonic crystal waveguide. It shows a detail of the band diagrams of a W1 photonic crystal slab waveguide in SOI, with a hexagonal lattice with $500nm$ period, and a bulk hole size of $300nm$. In the left figure, there are no proximity effects at the border holes. The middle and the right figure show the band diagram of the same structure, but with the border hole size $\oslash_{border}$ increased by $10nm$ and $20nm$ respectively. Even for such a small change, the characteristics of the guided mode change considerably. This can be easily noticed by the position of the mini-stop band (MSB) between the vertically odd and even mode. For

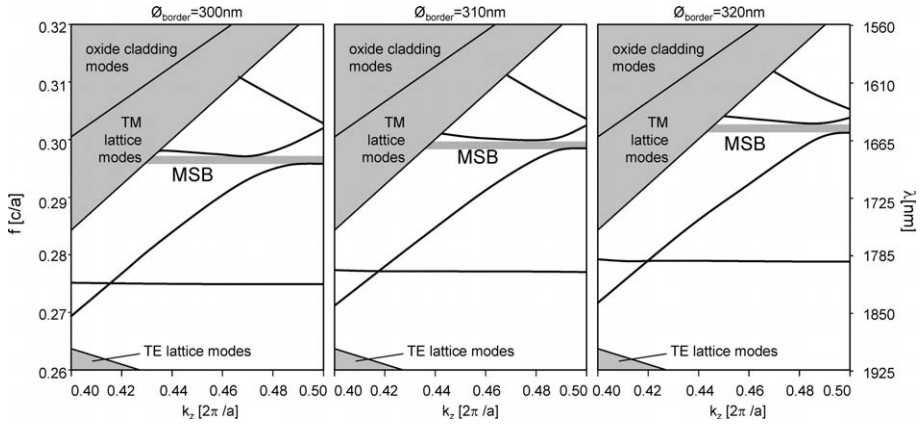

Figure 5.23: Impact of OPE on the band diagram of a W1 waveguide: The increase in the diameter of the border hole $\oslash_{border}$ causes a significant shift of the guided modes, which can be easily observed in the position of the mini-stop-band.

a change of $10nm$ in the border hole, the MSB centre wavelength shifts by almost $20nm$.

Note that proximity effects are not only a problem of optical lithography. With e-beam lithography, scattering electrons will also cause proximity effects in closely packed structures. However, because e-beam lithography is a serial writing process, the proximity effects just add up, and are therefore easier to model [99]. With optical lithography and deep UV lithography, the effect is coherent, and harder to predict.

As with the line-hole bias, to correct for OPE, the features on the mask should be altered. This is illustrated in figure 5.24. Holes near a lattice defect are printed smaller, and are therefore enlarged on the mask. It is evident that a good understanding of the OPE is necessary to design the structures with optical proximity correction (OPC).

To characterise optical proximity effects and the needed corrections in photonic crystals, we have included structures consisting of various photonic crystal waveguides along with a large number of bends, cavities and other possible components on the PICCO_1 mask. We then repeated this structure on a mask with many variations of bulk hole sizes, and corrections on corners and borders. This makes it possible for us to measure the optical proximity effects and the required corrections directly. Figure 5.25 shows the optical proximity effects for the holes in a $60°$ bend in a W1 waveguide. Such graphs are an interpolation from the OPE measured on our test structures and allow us to

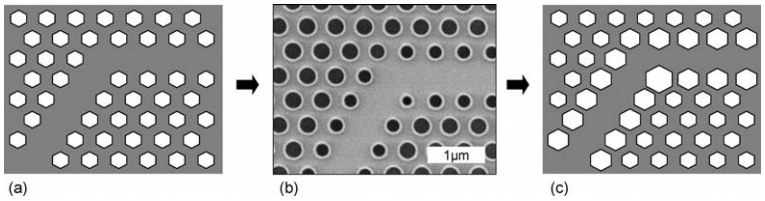

Figure 5.24: Principle of optical proximity corrections for photonic crystal structures. When uncorrected, the design on the mask (a) will print the defect holes differently than the bulk holes (b) due to optical proximity effects. To correct this, the design on the mask should be altered in advance (c).

apply the necessary OPC on future mask designs. In practice, this is done numerically instead of visually on the graph. For new structures, similar data can be measured on our test structures as the need arises.

This technique of establishing the necessary OPC limits us to structures similar to the which have been fabricated before. OPE can also be modelled, but due to the characteristics of the chemically-amplified photoresist this is not always reliable.

Not only holes are subject to optical proximity effects. Figure 5.26 shows a situation where the line width changes when the surroundings of a line is changed. We can see that the line width for the double line is larger than the line width for the dense lines. Similarly, the line width of the isolated line is larger than that of the double line. This is an example of a negative optical proximity effect, as intuitively, we would have expected the opposite behaviour: when the line is in a brighter environment, it should be printed narrower.

Figure 5.27 shows the bias between a double and an isolated line in four different combinations of line width and gap width. In all these cases, the OPE is negative, i.e. the isolated line is broader than the double line. We can see that the effect is not linear in the exposure dose. Therefore, where the line width or gap width needs to be corrected accurately, a detailed process characterisation is needed.

5.2.8 Mask Fabrication

Deep UV lithography is a mass fabrication process. Because the process requires large, high-quality photolithography masks, the design-to-device cycle is significantly longer than for typical research-oriented fabrication techniques, like e-beam lithography. The typical delays en-

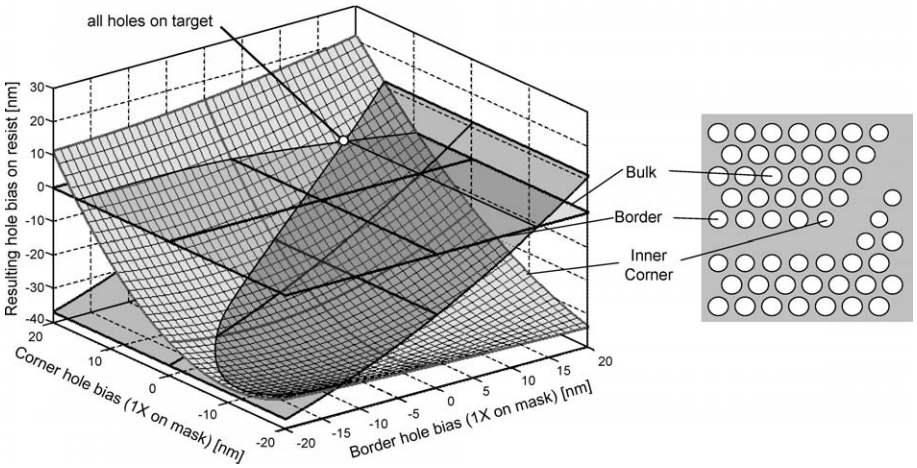

Figure 5.25: Optical proximity effects for a hexagonal lattice of holes with a pitch of 500nm. Design size of the bulk holes is 300nm, overexposed to 320nm. The graph shows the hole size for the border and inner corner holes when a bias (Optical proximity correction) is applied on the mask.

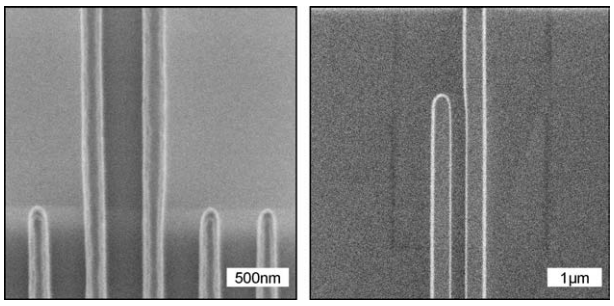

Figure 5.26: Optical proximity effects in line. Left: a double line and an array of lines. Right: An isolated line and a double line. In both cases, we can see that the line width increases when the line is more isolated. This is a negative proximity effect.

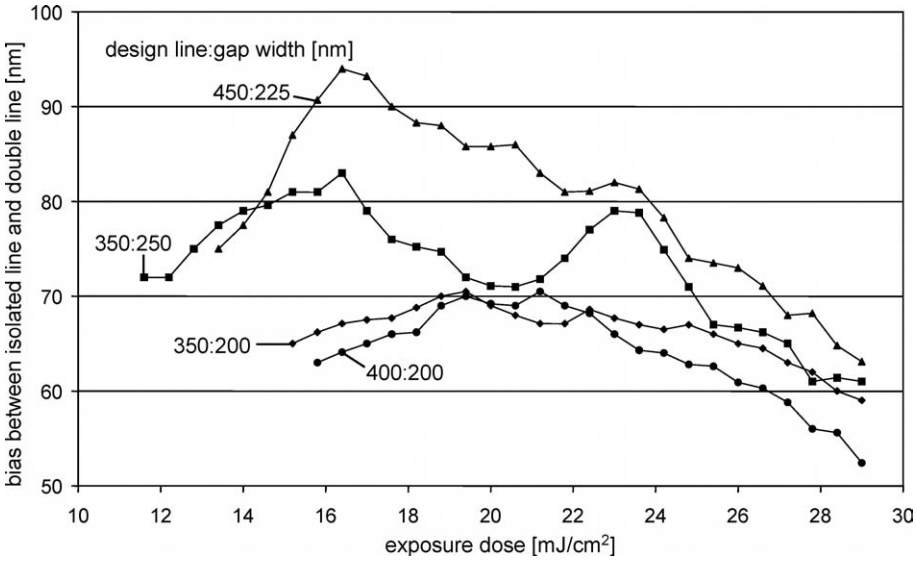

Figure 5.27: Optical proximity effects in lines: The bias between the line width of an isolated line and the line width of a double line is plotted as a function of exposure dose for 4 combinations of line width and gap width. Note that this is a negative proximity effect: the double line is narrower than the isolated line.

countered with deep UV lithography are mainly related with the mask-making:

- The design phase: In order to maximise the use of the expensive masks, a large number of designs should be incorporated. The simulation and layouting of these designs can take up several months.
- The mask fabrication: For standard CMOS masks, the typical fabrication time for a mask is up to 14 weeks, most of it taken up by data preparation on the side of the mask shop. For the first masks, extensive collaboration with the mask shop was needed to optimise the data of our non-standard nanophotonic structures.
- Mask validation and Line clearance: Once the mask is fabricated, it takes a couple of weeks before the mask is accepted in the standard manufacturing line.

Because the stepper operates with a reduction factor of $4X$, the mask features are 4 times larger on the mask than after lithography. This makes it possible to fabricate the mask with a very high accuracy. For the first mask, we chose a mask technology suitable for $0.18\mu m$ CMOS structures. Because the early experiments suggested that binary masks sufficed for most photonic wires and photonic crystals, no resolution enhancement techniques, like attenuated or alternating phase shifting, were required. In the course of this work, 4 such masks have been designed, designated PICCO_1 through PICCO_4. An overview of the mask layouts for these masks is given in appendix A

Like all high-end deep UV lithography reticles, the PICCO masks were defined with e-beam lithography. The pattern files for the individual cells were generated at INTEC with a number of scripts and converted to the standard GDS-II file format. These files were sent to IMEC in Leuven for verification. Also, IPC boxes were added to allow process validation for the mask fabrication and later for lithography and etching.

The actual mask fabrication was done by *Photronics* in Manchester, UK. Because these masks differ significantly from conventional CMOS masks, often several iterations were needed to make the mask on target. This, of course, yielded invaluable lessons, and with each subsequent mask the time from design to fabricated structures dropped, even though the designs on the masks grew ever more complex. Table

Step	PICCO_1	PICCO_2	PICCO_3	PICCO_4
Start design	01/03/2001	01/11/2001	01/11/2002	01/06/2003
Finish design	01/06/2001	15/03/2002	15/02/2003	31/07/2003
Mask Delivery	31/01/2002	01/09/2002	30/05/2003	01/10/2003
# iterations	4	2	2	1

Table 5.1: Development time for the four PICCO masks. For each subsequent mask, less design and fabrication time was needed, even though the designs grew in complexity.

5.1 gives an overview of the various masks and the time required to fabricate the masks.

Because of the steep learning curve, the first mask took 9 months from design to delivery. For the second mask, this time dropped to 6 months, while the third mask took only 4 months. At the time of writing, the PICCO_4 is being introduced in the IMEC clean room. Because of these large delays, it was sometimes needed that a new mask was ordered before the previous mask was delivered.

5.2.9 Resolution Enhancement Techniques

While we have demonstrated that the standard deep UV lithography process has sufficient resolution for both photonic crystals and photonic wires, there are some obvious reasons to increase the resolution even more. With higher resolution, smaller features are possible. For the larger structures, the process window will be larger, making the process more reproducible. Also, the optical proximity effects will decrease when the structures are further from the diffraction limit.

There are a number of techniques to increase the resolution of the lithography process, some more effective than others, and some only applicable to a specific type of structures:

- The most obvious way to increase the resolution is to move to a *shorter illumination wavelength*. As already discussed, this is a very expensive option. The current mainstream wavelength is $248nm$, with $193nm$ for high-end processes and $157nm$ under development.

- According to formula 5.6, we can also *increase the numerical aperture* of the projection system. A large NA need very complex optics, so there is a limit to the maximum NA. Current high-end steppers have an NA of over 0.8. Also, increasing the NA will reduce the depth of focus.

	λ	**NA**	**Resist**
A	$248nm$	0.63	UV3
B	$248nm$	0.7	UV3
C	$248nm$	0.7	TIS248
D	$193nm$	0.63	TIS193

Table 5.2: Lithography processes compared for the fabrication of nanophotonic structures.

- In order to allow for a small depth of focus, one can *use better resists*, which can be deposited in a thin layer and have a very nonlinear response.
- According to formula 5.6, we can also increase the resolution by *decreasing the spatial coherence factor* (increasing the fill factor σ). This is effective for periodic structures, but reduces the image contrast for isolated structures.
- Instead of changing the source size, one can change the shape. *Off-axis illumination* replaces the circular source by a dipole, a quadrupole or a ring. This, again, is beneficial only for periodic structures.
- *Phase Shifting Masks* (PSM) introduce a phase shift at the light-dark interface on the mask, effectively forcing the transmission to zero and enhancing the image contrast. However, PSM are costly compared to conventional binary masks.

Illumination Conditions

We studied the improvements that could be gained for nanophotonic devices by changing a number of exposure conditions: increasing the NA, using a thinner and better resist and switching to a wavelength of 193nm. We experimentally compared different processes to determine the influence of these factors. The different process parameters are listed in table 5.2.

As a comparison we determimed the largest 5% elliptical process window in the Focus-Exposure space for a triangular lattice of holes. This required measurements of hole diameters of the same structure at different exposure dose and focus. These measurements for $300nm$ holes with $500nm$ pitch are shown in figure 5.28. From these measurements, we can calculate the relative deviation from the target value. This gives us a 5% or 10% deviation contour. Within this contour, we

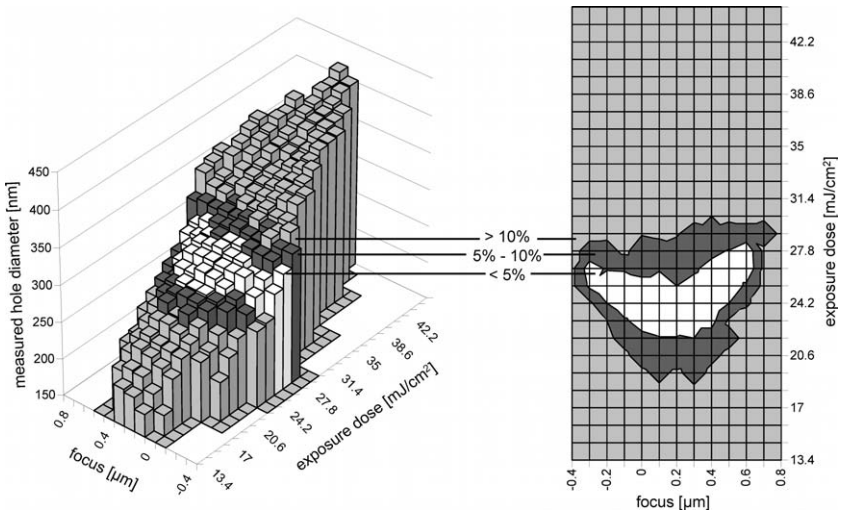

Figure 5.28: Measuring process windows. First, the hole diameter is measured for different focus and exposure dose (left). Then the absolute value of the deviation of the hole diameter relative to the target diameter is calculated and a 5% and 10% contour is plotted. Within this contours, the largest elliptical process window can be inscribed. The example shown is for a wavelength of $248nm$, $NA = 0.63$ and UV3 photoresist.

look for the largest elliptical process window. Because we have to compare different processes, we look at the relative exposure latitude, i.e. the absolute exposure latitude ΔE divided by the best exposure dose E_O.

We did this for the four processes in table 5.2 and holes with a pitch of $500nm$ and a target diameter of $300nm$, as well as for holes with a pitch of $400nm$ and a target diameter of $200nm$. For our reference process, we used a wavelength of $248nm$, an NA of 0.63 and the Shipley UV3 resist. We can see in figure 5.29 that we already have a rather large process window for the large holes. Figure 5.30 shows the process windows for the smaller periods, and we see that the process window is marginally sufficient. We then changed the NA to 0.7 with all other parameters the same. For the large holes, this increases both exposure latitude (EL) and depth of focus (DOF). For the smaller holes, increasing the NA meant a trade-off between EL and DOF, increasing the former and reducing the latter. For the third experiment we used an advanced bi-layer resist, which can be applied in a thinner layer and has a better nonlinear response. We kept the NA at 0.7 with a wavelength of $248nm$.

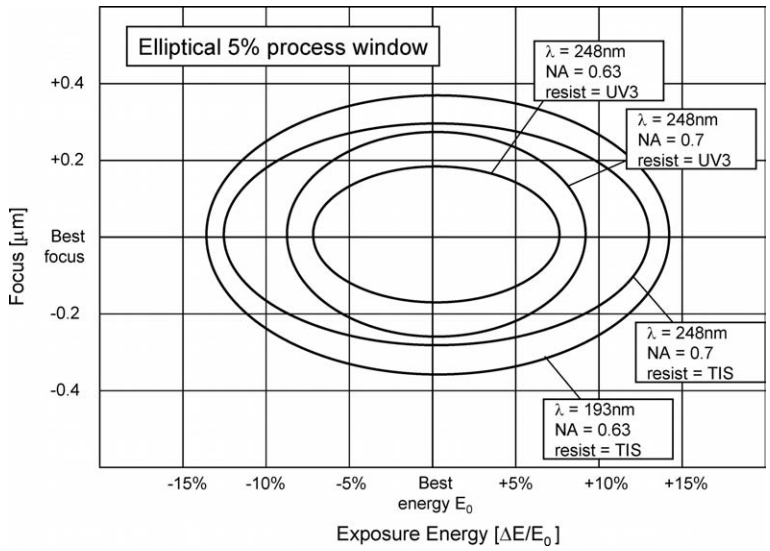

Figure 5.29: Process windows for $300nm$ holes with $500nm$ pitch, for 4 different lithography processes with different numerical aperture, illumination wavelength and resist.

Again, we see a significant improvement for both small and large holes. The last experiment uses a similar bi-layer resist, but for a wavelength of $193nm$ and an NA of 0.63. This enlarges the process window even more, with ample margin for both the large and the smaller periods.

We can see that by changing the process parameters, we can dramatically improve the resolution of the structures. Moreover, the improvements discussed here make no assumption of the type of structure and are effective for isolated and dense structures.

Off-axis Illumination

As we have seen on page 143, using a partially coherent source can increase the resolution for periodic structures. However, as we can see in figure 5.31, only part of the ± 1 diffraction orders are transmitted through the projection system. This has a negative effect on the image contrast. Part of the zeroth-order diffraction beam will not contribute to the image formation, as it has no counterpart in the first-order beams to interfere with. This part will just add useless light to the image and reduce the contrast.

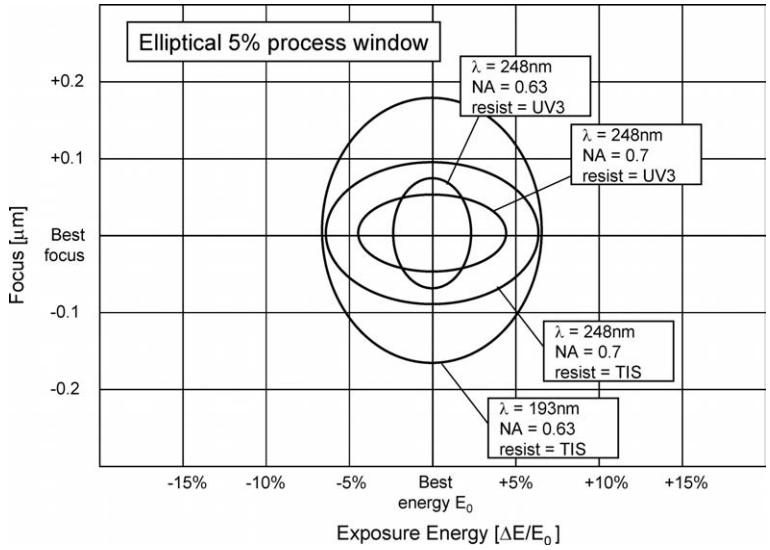

Figure 5.30: Process windows for $200nm$ holes with $400nm$ pitch, for 4 different lithography processes with different numerical aperture, illumination wavelength and resist.

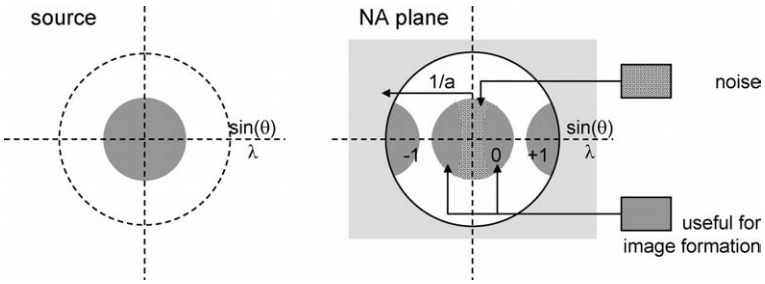

Figure 5.31: Noise from illumination with a partially coherent beam for small pitches a. For the diffraction orders ± 1 only part of the beam passed through the NA plane. The parts of the zeroth order beam that have no matching part in the higher order beams do not contribute to the imaging and therefore reduce the image contrast.

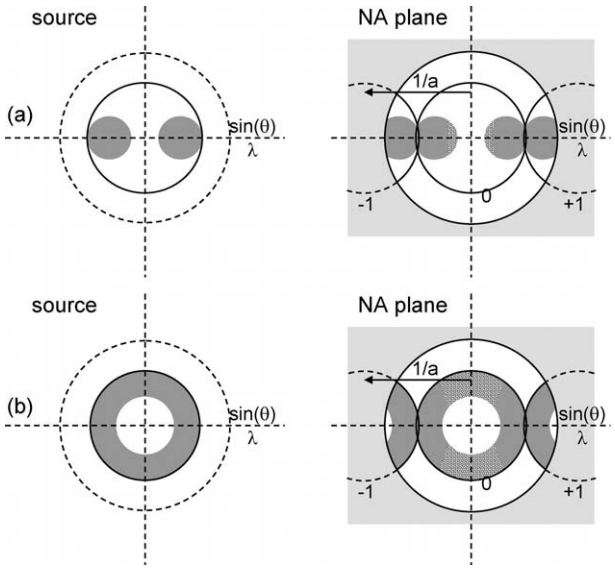

Figure 5.32: Off-axis illumination for contrast enhancement in dense periodic structures. (a) Dipole illumination: Because of the two lobes, a larger fraction of the zeroth order beam is used, resulting in less noise. This is only useful for periodic structures with a 1-D periodicity along the axis of the dipole. (b) Annular illumination. The principle is the same as for dipole illumination, but a ring-shaped source is used. This illumination geometry is useful for periodic structures in any direction.

To minimise this noise, the shape of the source can be optimised for periodic structures, as illustrated in figure 5.32. For 1-D periodic structures, *dipole illumination* can be used. The circular source is then replaced by a source with two lobes along the same axis as the periodicity. This way, the zeroth-order beam contains very little noise. However, this technique is only suitable for dense periodic structures in one direction. For square lattices, the technique can be expanded to *quadrupole illumination*, with a four-lobed beam. Because the largest intensity is not along the optical axis, these techniques are called *off-axis* illumination.

For periodic structures in a triangular lattice, one could use a six-lobed beam. Alternatively, *annular illumination* can be used, where the light source is ring-shaped. With this illumination technique, the image contrast is enhanced for periodic structures in all directions.

While off-axis illumination increases the image contrast for dense periodic structures, especially when the period is near the diffraction

limit, it degrades the image of isolated structures. Also, defects in periodic structures will not be correctly imaged, as optical proximity effects will be much larger. Therefore, these techniques are not suitable for the fabrication of nanophotonics.

Phase Shifting Masks (PSM)

Another area for possible improvement is the original image on the mask. We have already discussed optical proximity corrections, where the patterns on the mask is altered to obtain a correct image on the wafer. Typically, high-end deep UV masks are fabricated using e-beam lithography. Because the mask has a $4X$ reduction factor (the features on the mask are 4 times as large as they will be on the wafer), a mask can be made with very high accuracy. In appendix A the various mask layouts are described and the mask fabrication is given in more detail.

Up till now we have always assumed that the original image on the mask was a black-and-white pattern: light is either transmitted or not. In such *binary masks* there are only two possible intensities, and there is no phase information. For dense patterns, like photonic crystals, the images of the individual holes smear out and overlap, so the intensity in between the holes will not be zero. This is shown in figure 5.33a. For holes with a large fill factor, the resist in between the holes might be exposed beyond threshold. Similarly, for isolated lines, the images of the light area on both sides will partially expose the resist where the line should be.

This problem can be partially remedied by using a mask that contains phase information. In an *attenuated phase shift mask* (figure 5.33b), the dark area is not completely dark, but slightly transparent (typically 6%). Also, the light and dark areas have an opposite phase. This way, at the light-dark interface, the electric field E goes through zero. As a result, the intensity, which is proportional to E^2 is also forced to zero. This enhances the image contrast and better delineates the features.

Another application of phase shifting masks (PSM) is illustrated in figure 5.33c. In an *alternating phase shift mask*, the light areas consist of a periodic structure in which the phase alternates by π. This way, the profile of E is a sine curve with equal positive and negative peaks. The resulting intensity profile however, is also a sine, but with twice the spacial frequency of the original modulation. With this technique, very dense periodic structures can be made that are not possible with binary masks.

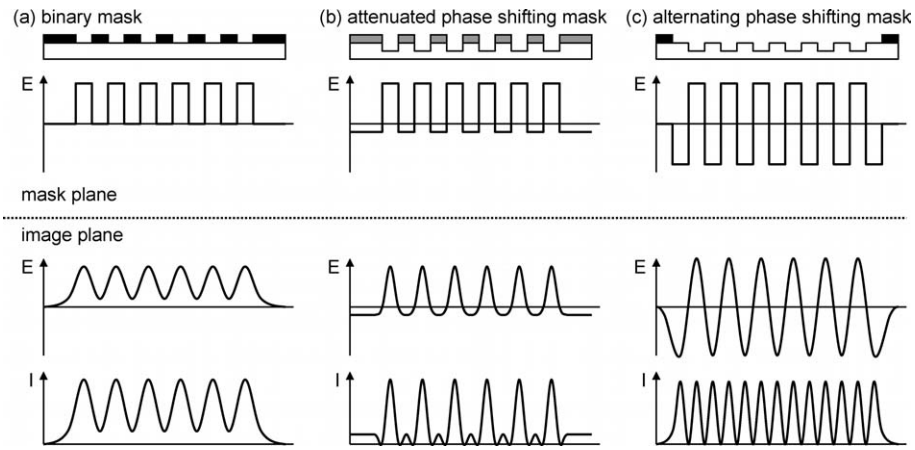

Figure 5.33: Phase shifting masks (PSM). (a) Binary mask. No phase shifting is used. For small and dense structures, the image is smeared out, resulting in badly defined light-dark interfaces. (b) Attenuated phase shifting. The dark area is not completely dark, but slightly transparent. The bright regions are shifted π with respect to the dark areas, so the intensity is forced to zero at a light-dark interface. (c) Alternating phase shifting. The periodic area is transparent, but with phase variations of π. The resulting intensity has half the period as the modulation of the PSM, allowing to make periodic structures with very small periods.

If we look at alternating PSMs from a viewpoint of the projection system, there is no violation of the diffraction limit. For dense gratings, only the 0 and ± 1 diffraction orders are passed through the lens, just as would be the case with a binary mask. However, with an alternating PSM, the zeroth order carries no power, and the orders $+1$ and -1 interfere directly, resulting in the frequency-doubled intensity pattern.

Alternating phase shifting seems like an attractive option for photonic crystals, as these are densely packed periodic structures. However, for a triangular lattice this is more difficult, as multiple phases are needed. Also, while it is possible to force the intensity to zero in the nodes between 3 holes, this is not possible in the thin line between 2 holes, where the resist is most likely to break down. Also, near intentional defects, alternating phase shifting is very difficult to use.

Phase shifting masks offer a serious resolution enhancement. While alternating PSMs can be used for dense periodic structures with a small period, they have a limited applicability. On the other hand, attenuated PSMs can easily be used for most structures to enhance the image contrast. The drawback of using this advanced mask technology is the additional cost, as phase shifting masks can carry up to three times the price tag of a conventional binary mask.

5.2.10 Conclusion

We have demonstrated that deep UV lithography, in particular at a wavelength of $248nm$, is capable of imaging nanophotonic components such as photonic crystals and photonic wires. While, logically, $193nm$ offers better resolution, we have shown that for most purposes $248nm$ lithography has a sufficiently large process window. However, a detailed process characterisation was needed to compensate for line-hole bias and optical proximity effects, especially in densely packed structures like photonic crystals. For nanophotonic structures operating at shorter wavelengths, like $1.3\mu m$, or some critical wavelength-selective elements, additional resolution could be required and justify the higher cost of $193nm$ lithography.

5.3 Etching

Following lithography, the structures defined in photoresist should be transferred to the underlying SOI substrate. As we have seen, keeping the refractive index of the bottom cladding as low as possible by

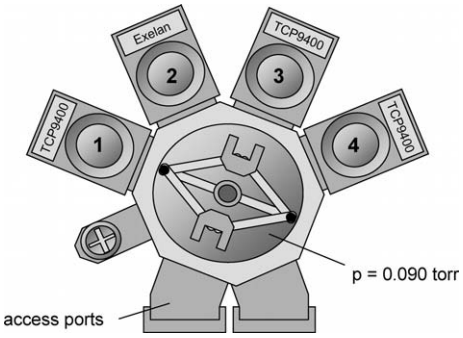

Figure 5.34: top-view schematic LAM_A6 platform with 4 process modules. PM3 is used for Silicon etch, while PM2 is used for the oxide etch. Wafers can be moved between process modules without exposure to the external atmosphere.

etching the holes deep into the buried oxide yields better optical properties. Another rationale for deeply etched holes is that the light in the top Silicon layer will hardly feel the bottom of the holes, and won't be scattered. However, as we will show, this deep etching causes substantial sidewall roughness. After trying several techniques for roughness reduction, we found that the best results could be achieved by not etching the oxide altogether.

5.3.1 First Experiments: Thin Oxide

For the first experiments, we used SOI wafers with a top Silicon layer of $205nm$ and an oxide of $400nm$. The etching of this structure is performed in two steps: first the top Silicon layer is etched, then the underlying oxide. The equipment used for etching is a LAM_A6 platform with four process modules (PM). Figure 5.34 shows the platform with three *TCP9400* modules and an *Exelan* module. For the Silicon etch, PM3 is used with a low-pressure/high-density $Cl_2/O_2/He/HBr$ chemistry. The etch recipe consists of a break-through etch and a main etch, with different chemistry.

After the Silicon etch, the wafer is transferred to the Exelan module (PM2) without exposure to the atmosphere. The oxide etch is done with a medium density CF_4/O_2-plasma chemistry at medium pressure. The etch depth can be controlled with the etch time.

Because of the limited oxide thickness, in the second step we can completely etch through the oxide cladding without consuming the re-

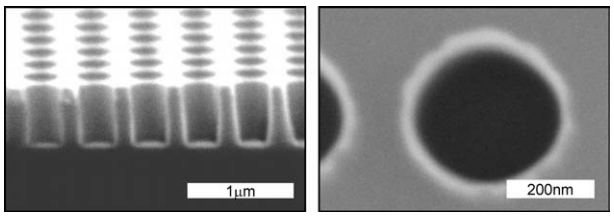

Figure 5.35: Etch results on SOI wafers with $400nm$ oxide. Left: The cross section shows smooth sidewalls. Right: In top-down, there is also little roughness.

sist. Figure 5.35 shows a cross section and a top-down view of a etched holes. In both views we see that the sidewalls are relatively smooth, even with the deep etch.

However, as we have seen in chapter 4, and as will become clear from the measurements in the next chapter, the $400nm$ oxide layer is too thin and causes high substrate leakage.

5.3.2 Deep Etching

To reduce the substrate leakage, we switched to wafers with a thicker oxide buffer. In order to eliminate scattering at the bottom of the etched holes, they should be etched as deeply as possible into the oxide.

The development of a deep etch process was less straightforward than originally expected. Because the resist is used directly as an etch mask, only a limited etch depth is possible before the remaining resist breaks down. Also, such deep etch processes for small features in a multi-layer substrate are seldom required in CMOS devices. It therefore took several months to get an acceptable etch process. Figure 5.36 shows a cross section of etched holes. One significant side effect of the deep etching is the appearance of sidewall roughness. For deeply etched structures, as the photonic wires illustrated in Figure 5.37, the sidewall becomes very irregular. Because sidewall roughness causes scattering of light as it propagates through the waveguide, this has to be avoided.

5.3.3 Roughness Reduction

Rough surfaces cause scattering. As discussed in chapter 4, the amount of scattering is approximately proportional to the square of the difference in dielectric constant on the surface interface, so the effect is dra-

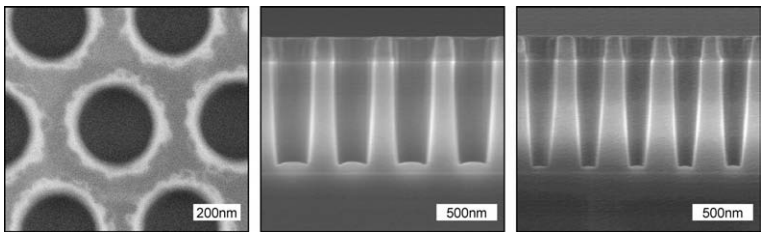

Figure 5.36: Etch results on SOI wafers with $1\mu m$ oxide. Left/middle: $500nm$ period / $300nm$ diameter. Right: $400nm$ pitch / $240nm$ diameter.

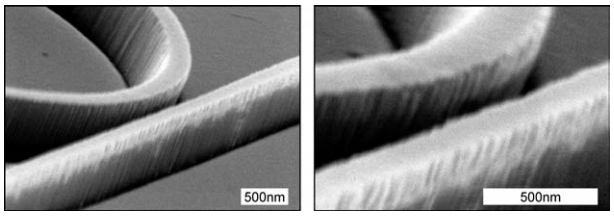

Figure 5.37: Details of a deeply-etched photonic wire with ring resonator. The blow-up on the right clearly shows the vertical roughness caused by the etching process. The straight wire is $400nm$ wide, the ring has a width of $500nm$.

matic for high-contrast material systems, like Silicon-on-insulator. To reduce these scattering losses, one can either avoid creating rough surfaces, or smoothen existing surfaces. We have examined a number of possibilities to reduce the sidewall roughness.

Before etching, one can smoothen the resist with a plasma treatment, and eliminate the possible irregularities that are already present in the photoresist. However, the most obvious approach is to abandon deep etching altogether, and only etch the top Silicon layer. We have also experimented with the deposition of additional oxide on top of this structure, in order to improve the vertical symmetry. An alternative approach is the partial oxidation of the top silicon layer to smoothen the roughness.

Resist Hardening and Trimming

In order to avoid creating rough surfaces, we have explored two options. One is to expose the photoresist to a plasma treatment before etching. This technique is especially used for the smoothening and narrowing of CMOS gates. However, our experiments showed that this could only be done at the expense of resist thickness, and loss of line

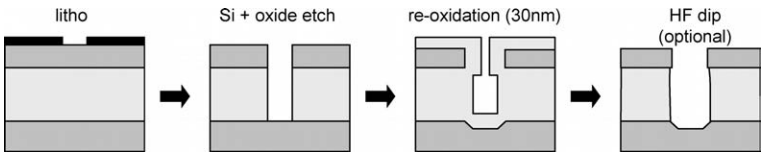

Figure 5.38: Roughness reduction using thermal oxidation of the Silicon sidewalls. The oxide growth is at the cost of the top Silicon layer.

width. Because a thick layer of resist is needed for deep etching, this was not an option.

Oxidation

We also tried to alleviate the sidewall roughness by treating the deeply etched structures. It has already been shown that thermal oxidation of Silicon-on-insulator waveguides can smoothen the sidewalls of both photonic crystal waveguides [100] and photonic wires [42, 101]. The principle is illustrated in figure 5.38. As the rate of oxidation is well documented and can be controlled quite accurately with the temperature, this is a reliable way of reducing roughness.

In figure 5.39 we see an example of a photonic crystal hole with different amounts of oxidation. The first impression is that the oxidation increases the amount of roughness on the sidewalls, blowing up the existing irregularities. However, because the new roughness is on an oxide-air interface instead of on a Silicon-air interface, the impact is less dramatic. Although this is hard to establish experimentally, we can assume that the underlying Silicon-oxide interface is smoother, due to the diffuse nature of the oxidation process. We can also see that the volume of the top layer increases after oxidation, creating a rounded core layer in cross section.

While oxidation can improve the roughness in the top Silicon layer, it has little or no effect on the underlying oxide layer. Therefore, light will still be scattered by the roughness in the cladding. A solution to that problem would be to remove the cladding layer, creating a free-standing membrane. While this is a valid option for small areas of photonic crystal structures, it is harder to achieve for photonic wires, as the line waveguides are unsupported by the substrate.

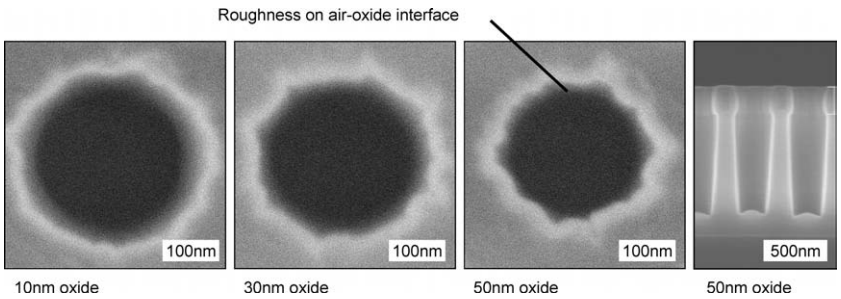

Figure 5.39: Deeply-etched photonic crystal holes after partial oxidation of the top silicon layer. Although the sidewall roughness seems more severe after oxidation, the irregularities are on the oxide-air interface, causing far less scattering than on the silicon-air interface.

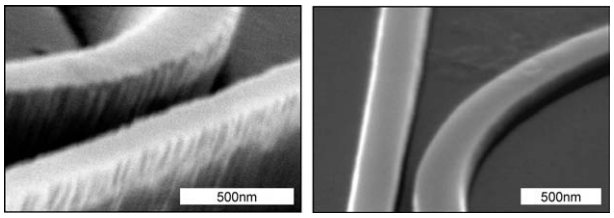

Figure 5.40: Comparison between deep and shallow etching. The SEMs show a detail of the coupling section between a straight photonic wire and a ring resonator. Left: Deep etch through both silicon and oxide. Right: Silicon-only etch. It is clear that the deep etching causes significantly more sidewall roughness.

5.3.4 Silicon-only Etching

While from a photonics point of view the oxide etching is attractive, it proved very difficult to keep the sidewall roughness acceptably small. Therefore, we also experimented with etching only the Silicon layer. By leaving out the second etch step, the sidewall roughness was drastically reduced. This can be seen in figure 5.40 and figure 5.41. The sidewall roughness of the photonic wire is almost completely eliminated.

Resist Hardening

The etch process gives rise to a considerable bias between lithography and etch, i.e. the etched holes were typically $50nm$ to $70nm$ smaller in diameter than the holes after lithography, as can be observed in figure 5.41. This can be compensated lithographically for isolated structures,

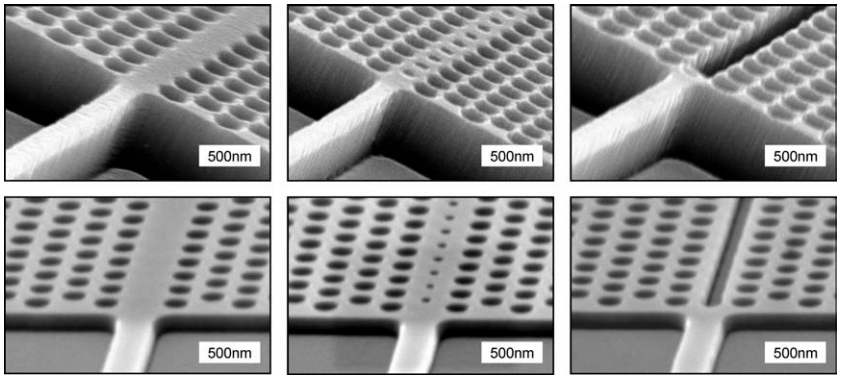

Figure 5.41: Photonic crystal waveguides fabricated with deep UV lithography and dry etching. Top row: 3 waveguide designs fabricated with deep etching, through both top silicon layer and oxide. Bottom row: The same designs but fabricated with only Silicon etch. Although the sidewalls are much smoother in the latter structures, the bias between lithography and etch makes the holes too small.

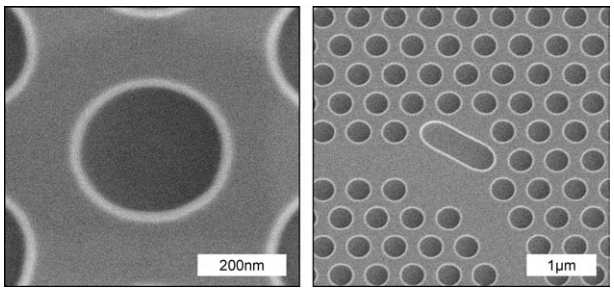

Figure 5.42: Top-down SEM pictures of structures fabricated with a resist-hardening + Silicon-only etch process. The sidewalls are very smooth.

but with the densely packed photonic crystal holes, the bias is too large to compensate with overexposure. This is because the narrow resist lines between the holes break down when the holes are overexposed beyond a certain radius.

Experiments with a resist-hardening plasma treatment showed that we could successfully reduce the bias to less than $30nm$. This is sufficiently small to be compensated by overexposure during lithography. Top-down inspection suggests very smooth sidewalls, as illustrated in figure 5.42.

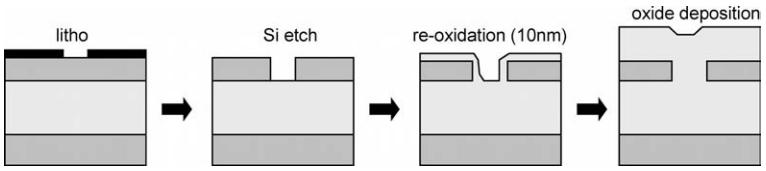

Figure 5.43: Roughness reduction using a short thermal oxidation and making the structure symmetric using oxide deposition.

Oxide Deposition

While etching only the silicon can reduce the roughness, we end up with an asymmetric layer structure. However, we can make the layer structure symmetric again by adding a top cladding with the same refractive index as the underlying oxide. This is illustrated in figure 5.43: After a silicon-only etch, which causes little sidewall roughness of its own, we can do a short thermal oxidation. After that, we can do an oxide deposition, which makes the structure symmetric in the vertical direction.

We used a *chemical vapour deposition* technique with a chemistry based on $SiO_4(C_2H_5)_5$ (TEOS). With an optimised process, even deep and narrow holes can be filled without creating voids.

A first experimental example of this technique is shown in figure 5.44. While the Silicon is not fully etched, we can see that the oxide deposition creates a smooth, planar top cladding and no artefacts (like voids) in the photonic crystal holes. One side effect of this technique is that the SOI structure is sealed from the outside world. While this can be advantageous for commercial components, for research purposes it makes close inspection of the structures with an SEM impossible. While this technique is very promising, we have not made actual devices with it.

5.3.5 Conclusion

Developing a good etching process for nanophotonic structures in SOI has proven to be more of a challenge than anticipated. The deep SOI etching, both into Silicon and oxide, leaves the sidewalls with a considerable amount of roughness, which is unacceptable for photonic waveguides. Experiments to reduce the roughness by thermal oxidation of the top silicon layer show that the roughness can be reduced, but not eliminated. The alternative, etching only the top Silicon layer, per-

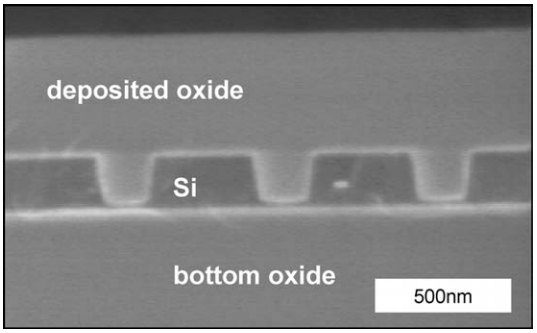

Figure 5.44: Cross section of photonic crystal holes with $500nm$ pitch after $5nm$ oxidation and oxide deposition. Note that there are no voids in the deposited oxide, and that the top surface is planarised.

forms much better with respect to sidewall roughness. However, the etch process introduces a bias with respect to the lithography which is too large to easily compensate for with overexposure. In combination with a plasma treatment of the resist after development this bias can be eliminated. An additional advantage of the Silicon-only etch is that it allows for subsequent deposition of oxide, effectively creating a vertically symmetric structure.

5.4 The Final Fabrication Process

The result of this comparison between various processes is a process suitable for a wide range of nanophotonic components. For photonic wires, we have demonstrated wire widths down to $300nm$. For photonic crystals, the limitation is the lattice pitch. For triangular lattices with a period as small as $400nm$ we still have a marginally acceptable process window.

The process flow is illustrated in figure 5.45 [51]. For the lithography we use the standard $248nm$ process, with $800nm$ of UV3 resist, a spatial coherence factor $\sigma = 0.4$ and a numerical aperture $NA = 0.63$. For etching, we limit ourselves to the top silicon layer, but with resist hardening prior to etching to compensate the litho-etch bias. While additional roughness reduction should be possible with oxidation, and deposition can be used to make the structure vertically symmetric, we have not included these steps for device fabrication.

Even with optimised designs, not all structures will have the same dose-to-target. Therefore, all devices are fabricated with a variety of ex-

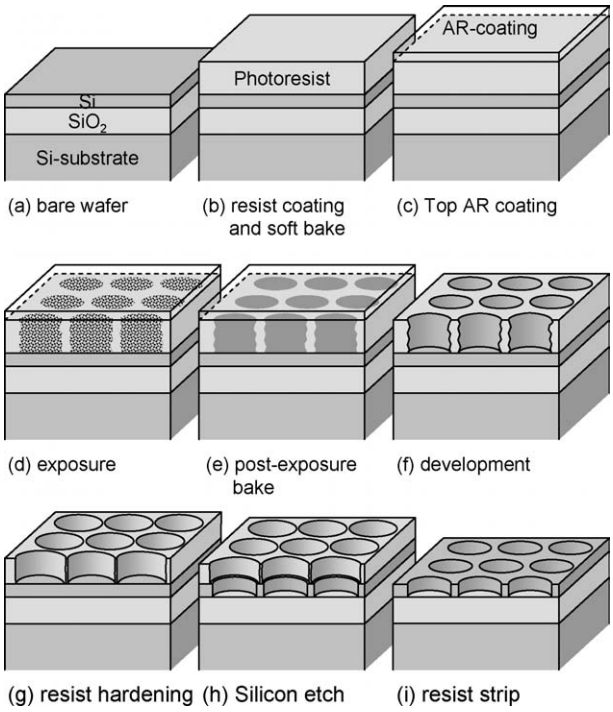

Figure 5.45: The optimised process flow for the fabrication of nanophotonic structures in SOI.

posure doses. As the dose can be changed from die to die on a wafer, a linear scan of the dose is done in the horizontal direction. In the vertical direction, all dies in a column are identical. This way, we have a large variety of structures on a single wafer, while still keeping multiple dies with the same structure parameters.

5.5 Post Processing

To characterise the fabricated structures, light has to be coupled into the waveguides. The most straightforward way of doing this, as will be illustrated in chapter 6, is by focusing a spot from a fibre into the waveguide at a cleaved facet. Similarly, the transmitted light has to be coupled out of the structure at a cleaved facet to be collected onto a power detector. In our designs, the total length of the structures is typically $5 - 7mm$, so after cleaving, the incoupling and outcoupling facets are typically $4 - 6mm$ apart.

Cleaving a nice facet in Silicon is not straightforward. Unlike III-V semiconductors, Silicon is more resilient to stress. With a wafer thickness of $725\mu m$, cleaving two smooth facets $4mm$ apart is very difficult to do reproducibly. Therefore, the thick substrate has to be thinned by backlapping.

The procedure is illustrated in figure 5.46. As the lapping equipment at INTEC does not support $200mm$ wafers, the wafer has to be divided into dies of approximately $4 \times 4cm^2$. Typically, we make sure to have at least one sample for each exposure dose. To protect the structures on the die, a layer of photoresist is deposited and baked. The die is then glued onto a glass plate with bee's wax heated to $130C$. After cooling, the glass plate is mounted on the backlapping machine with a vacuum chuck. Depending on the size of the die, the backlapping with a steel plate and a suspension of alumina powder can take 1 to 4 hours. The $725\mu m$ substrate is thinned down to about $250\mu m$.

To remove the thinned sample from the plate, the wax is heated again. The wax residue on the sample is dissolved in TCE (C_2HCl_3). The TCE and the photoresist are then removed with acetone. Finally, the sample is cleaned with isopropanol (IPA) and DI water. To cleave the clean sample, scratches are made on the top with a diamond needle and pressure is applied on the bottom.

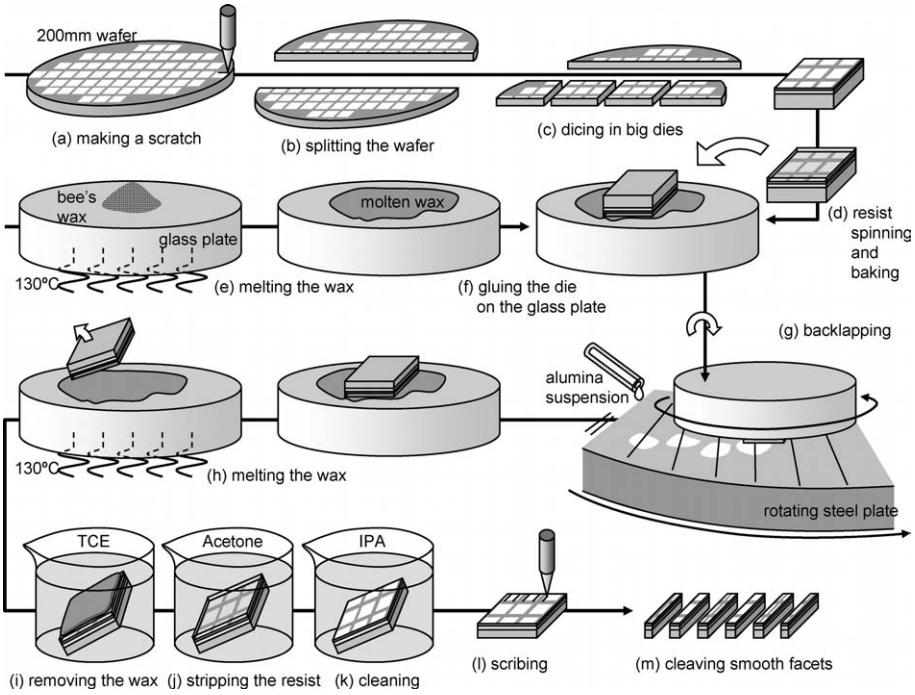

Figure 5.46: Postprocessing to make incoupling facets. The $200mm$ wafer is split in dies of approximately $4 \times 4cm^2$ (a,b,c). These dies are coated with resist (d) to protect the structures and glued upside-down (f) on a glass plate with hot bee's wax (e). Then, the $725\mu m$ thick substrate is backlapped down to about $250\mu m$ (g). The thin sample is removed from the plate by melting the wax again (h), and the residual wax is dissolved (i) in C_2HCl_3 (TCE). Then, the resist coating is stripped with acetone (j) and the sample is cleaned with isopropanol (IPA) and DI water (k). With a diamond needle, scratches are made where the incoupling facets should come (l) and the sample is cleaved by applying pressure on the bottom (m).

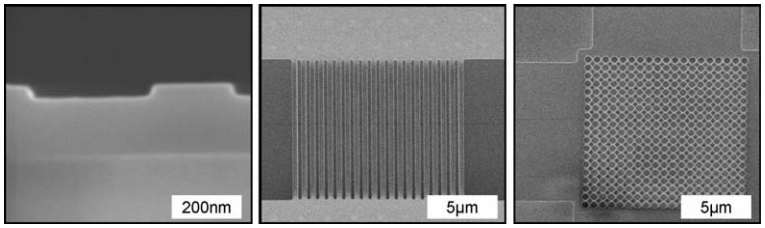

Figure 5.47: Fibre couplers in SOI fabricated with deep UV lithography. Left to right: Cross section of a 1-D fibre coupler, top-down of a 1-D fibre coupler, top-down view of a 2-D fibre coupler/polarisation splitter.

5.6 Fibre Couplers for In and Outcoupling

The post-processing cycle can take up to three days for one wafer, and then just for one or two samples at each exposure dose. Because this is a labour-intensive and time-consuming process, an alternative incoupling strategy was explored where there is no need for cleaved facets.

5.6.1 Fibre Couplers

The vertical fibre couplers introduced at the end of chapter 3 provide a way of coupling light into a broad waveguide from a vertically oriented fibre. This way, measurements can be done directly on the wafer without any post-processing. Therefore, we combined fibre couplers with nanophotonic waveguides[4].

Unlike photonic crystals and photonic wires, the fibre couplers are not etched completely through the top Silicon layer. Depending on the type of grating and the fill factor, the optimal etch depth is between $40nm$ and $80nm$. Figure 5.47 shows a cross section of an etched fibre coupler and a top-down micrograph of a 1-D and a 2-D grating. As these structures were fabricated in a single processing step, the trenches that define the broad waveguide have the same etch depth as the fibre coupler. Therefore, to use the fibre couplers together with deeply etched nanophotonic waveguides, the gratings should be fabricated in a different process step than the waveguides.

The fibre couplers provide an easy means to couple into and characterise nanophotonic structures. However, where a large bandwidth is needed ($> 60nm$), the fibre couplers can't couple enough power

[4]The simulation, design and measurement work on the grating-based fibre couplers has been carried out by Dirk Taillaert in the context of his PhD.

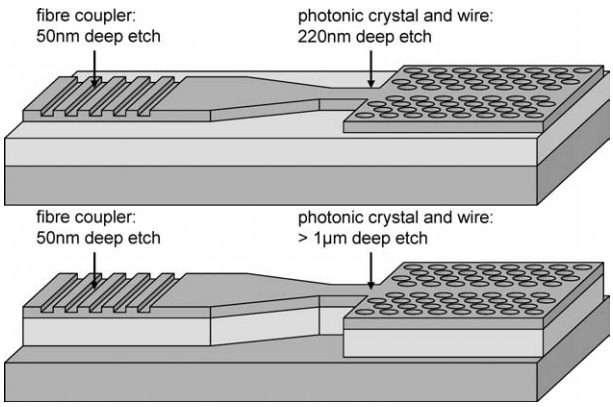

Figure 5.48: Two-step processing. The fibre couplers need another etch depth than the nanophotonic waveguides. Top: Nanophotonic waveguides with a Silicon-only etch. Bottom: Nanophotonic waveguides with deep etching.

over the complete wavelength range. Therefore, in some cases, cleaved facets are still required.

5.6.2 Two-step Processing

While photonic crystals and photonic wires require deep etching, at least through the top silicon layer, the fibre couplers need only about $50nm$ etch depth [34]. Figure 5.48 shows the situation with both the Silicon-only etch and the deep etching. The two different etch processes mean that the fibre couplers and the waveguides should be fabricated in a separate run. Because deeply etched structures can cause a problem for the lithography (as the wafer will have too much topography and the resist layer will not have a homogeneous thickness) the shallowly etched fibre couplers should be fabricated first. In order to align both types of structures, we can rely on the alignment tools available in the stepper. Instead of providing alignment features on the individual dies, the stepper takes care that the dies are well aligned to 6 alignment markers on the wafer. When the stepper knows the absolute position of the different structures on the reticle, it can carry out the necessary alignment with an accuracy of $50nm$ or better. This is much faster than die-based alignment, as the alignment needs to be done only once per wafer.

We have incorporated such two-step structures on the PICCO_3 and PICCO_4 reticles. By standardising the locations of the fibre couplers,

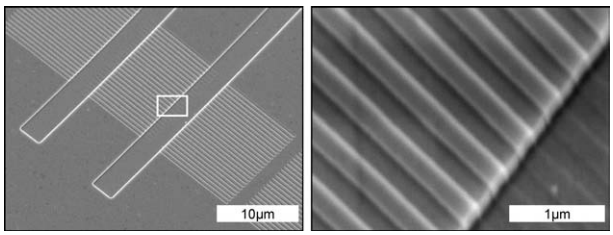

Figure 5.49: Fibre couplers fabricated with two-step processing.

we can reuse a single mask cell with fibre couplers for all other structures that adhere to the standard positions of the broad incoupling waveguides. The PICCO_3 mask contains such 1-D fibre couplers. A fabricated example is given in figure 5.49. In the detail we can see that the fibre coupler grating, which has been etched first, is still imprinted in the trenches of the broad waveguide, which has been etched afterwards.

Next to the 1-D couplers, we have also included an experimental polarisation-diversity circuit with a 2-D fibre coupler/polarisation splitter on the PICCO_4 mask.

5.7 Conclusion

We have demonstrated the use of CMOS technology, and more specific deep UV lithography for the fabrication of nanophotonic structures. While deep UV lithography is already the workhorse for CMOS fabrication, nanophotonics impose a completely different set of requirements on the fabrication process. As alignment issues are much more critical a single litho step should be used to align all structures correctly. However, this can introduce a considerable mismatch between the different types of components because of the different dose-to-target for holes, lines and other structures, as the optimal lithography conditions for each type of structure are different. Therefore, detailed process characterization is required to determine the correct bias for each individual structure. Also, the dense nature of photonic crystals gives rise to optical proximity effects. We have fabricated test structures to experimentally measure optical proximity effects, and the necessary corrections to apply on the mask.

For the fabrication we used photoresist as the etch mask. Because the deep etching, through both the silicon and the oxide layer, caused

a large amount of sidewall roughness, we had to look for roughness reduction techniques. The most promising proved to be not to etch the oxide, but only the top Silicon layer. Optionally, an additional oxide layer can be deposited on top of the silicon to make the structure more vertically symmetric.

In order to avoid the time-consuming wafer thinning, we can include grating-based fibre couplers in a separate process step.

Summarised, we can state that deep UV lithography has the potential to provide a commercial mass-manufacturing base for nanophotonic components. While detailed process characterisation and development is required, we have shown that many fabrication hurdles can be overcome.

Chapter 6

Measurements

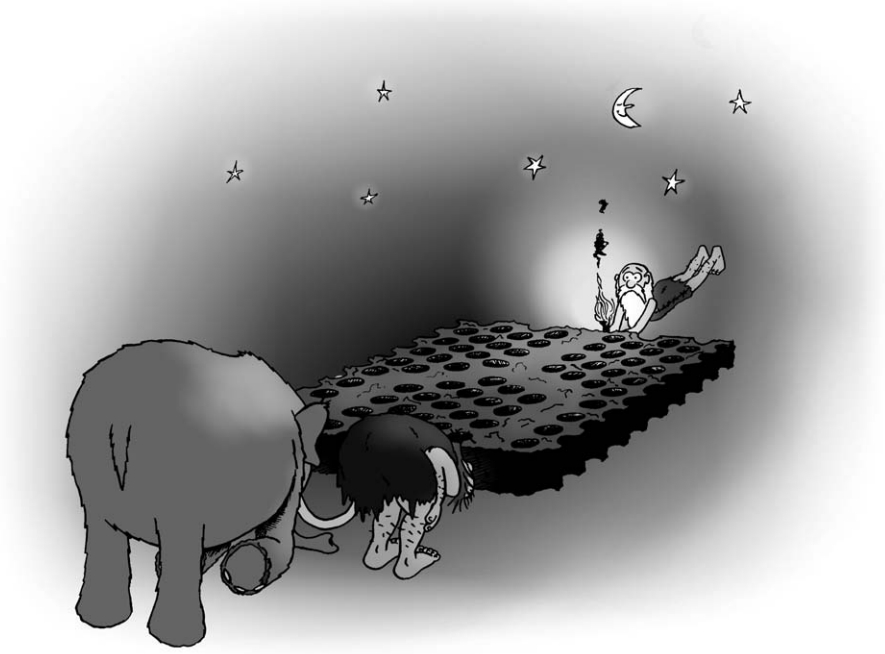

"No, not a spark coming through..."

Chapter 6

Measurements

Once nanophotonic components have been fabricated, they need to be characterised to determine their optical properties. In this chapter we will explain the various measurement techniques and a number of noteworthy results obtained with these measurements. As both the theoretical knowledge of nanophotonic components and the technological expertise with the fabrication process grew, the designed components have become increasingly better.

6.1 Measurement Setup

To characterise the optical properties of the nanophotonic components, we measure the transmission of light through the component as a function of wavelength. For this purpose, an extensive measurement setup is available at the INTEC group. The light source is a *Photonetics* tunable laser with a wavelength range from $1500nm$ to $1640nm$ and an output power of at least $2mW$ in this range. The laser can be controlled from a PC through a GPIB cable.

For this work, we have used two measurement techniques. With *end-fire measurements*, we used cleaved facets for incoupling and outcoupling. This method is inherently simple, but required very accurate alignment. Also, because of the mode mismatch between the fibre and the facet, the incoupling efficiency is rather low.

Aa an alternative, we used *vertical fibre couplers* for both incoupling and outcoupling. While the alignment requirements are less strict compared to the end-fire method, the fibre couplers have only a limited bandwidth.

6.1.1 End-fire Measurement

For end-fire measurements, light is launched into optical waveguides through a cleaved facet and the transmitted light is collected at the opposite side of the sample. The setup for these measurements is illustrated in figure 6.1.

The light from the laser is coupled to a polarisation-maintaining fibre, which in turn is spliced to a lensed fibre for incoupling. Because the lensed fibre, which is not polarisation maintaining, can induce a polarisation conversion, an additional polarisation controller is needed which can change the fibre polarisation by applying torque on the fibre through three wheels.

The sample with nanophotonic waveguides is mounted on a 2-D translation stage (x and y direction) with submicron controls. The sample is held in place by a vacuum chuck. The lensed fibre is mounted on a 3-D translation stage. A coarse alignment is possible through a binocular microscope mounted above the sample. The spot of the lensed fibre is focused on a cleaved facet of a $3\mu m$ wide ridge waveguide. The waveguide is then tapered down to a useful component. The outcoupling is at the opposite facet through a similar waveguide, but the light is collected by an objective lens and focused onto a power detector. The objective lens is also mounted on a 3-D translation stage.

The power detector, too, is connected to the PC through GPIB so automatic wavelength sweeps are possible. In order to avoid stray light from falling onto the power detector, an adjustable aperture is placed in front of the detector. The detector is placed on a pivot so it can swing away from the optical axis. The image of the outcoupling facet is then projected onto an infrared camera to facilitate the alignment of the setup and make an estimate of the intensity profile in the outcoupling waveguide.

In practice the end-fire technique has some considerable disadvantages. Firstly, the coupling efficiency from the lensed fibre to a thin SOI waveguide is of the order of $-10dB$, and the outcoupling to the detector adds approximately $-8dB$ loss. Moreover, because the alignment of the lensed fibre is very critical, the coupling efficiency can vary from one waveguide to the next or between samples. Also, the coupling efficiency depends very much on the facet quality. To produce clean facets, extensive post-processing is required, as already discussed in chapter 5.

As already said, we used a tunable laser to scan the wavelength-dependent behaviour of our structures. However, the facets create a

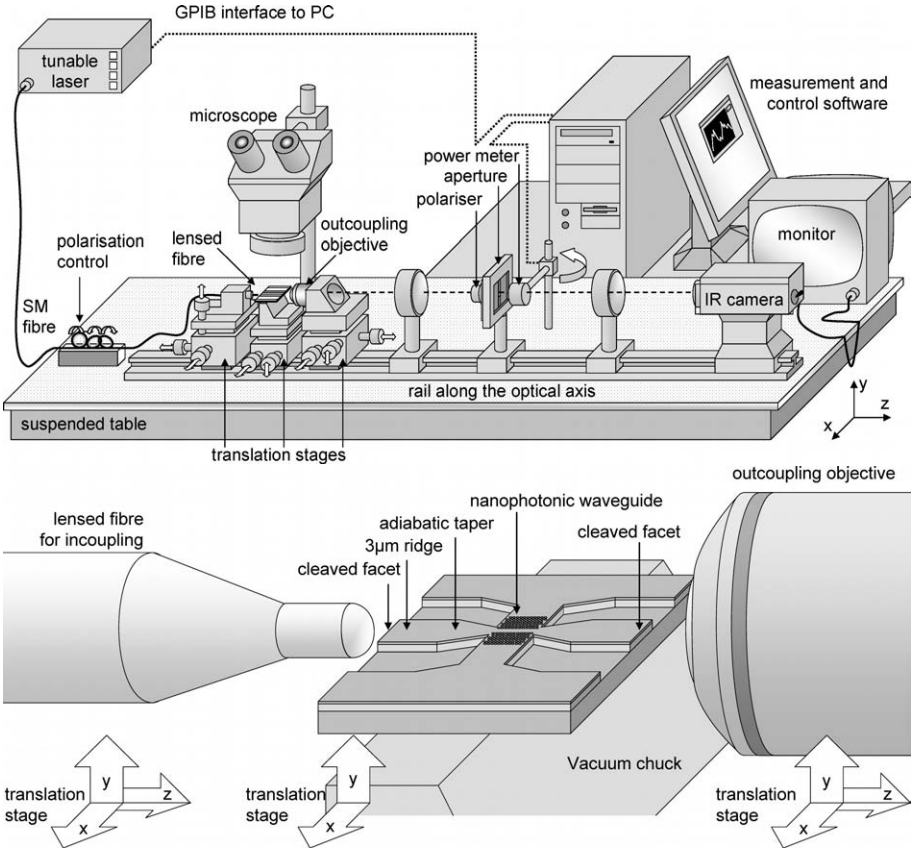

Figure 6.1: The measurement setup for end-fire measurements. Light from a tunable laser is coupled into a nanophotonic circuit by means of a lensed fibre. The polarisation of the light is controlled by applying torque to the fibre with a set of polarisation control wheels. The light transmitted through the sample is collected by an objective and projected on a power detector. The power detector can be pivoted away from the optical axis, so the image of the outcoupling facet can be observed directly by an infrared camera. Both the laser and the power detector are computer controlled through a GPIB interface. For alignment, the incoupling fibre, the sample and the outcoupling objective are placed on individually addressable translation stages.

cavity, which produces a strongly oscillating transmission as function of wavelength. As the laser has a fixed line width, the only way of determining the correct transmission curve is by scanning the wavelength range with a sufficiently fine resolution. Scanning the whole wavelength range in this way can take several hours, making this a slow measurement technique.

The alternative, using a broadband light source and a spectrum analyser with an adjustable wavelength window, would allow us to average out the oscillating signal caused by the Fabry-Perot cavity speeding up the measurements. However, for this work we had no broadband light source at our disposal.

6.1.2 Measuring with Fibre Couplers

When using vertical fibre couplers, much of these disadvantages are eliminated. The fibre couplers have a coupling efficiency of about 20%, which is better than with a lensed fibre. Also, the alignment tolerances are much better, and no post-processing of the samples is required.

Figure 6.2 shows the measurement setup with the fibre couplers. The light from the laser is coupled to a single-mode fibre which is positioned above the sample. The fibre coupler grating is intentionally optimised for oblique incoupling in order to avoid backreflection and the formation of a cavity inside the sample. Therefore, the fibres need to be tilted 10° with respect to the vertical axis. The polarisation of the light in the fibre is again optimised by a polarisation controller. In order to position the fibre correctly, a microscope camera and a lamp are used. When the fibre is brought in proximity of the sample, a shadow is cast on the sample. From the relative position of the fibre and the shadow on the video monitor, the fibre can be positioned using a 3-D translation stage. This crude positioning is done for both input and output fibres.

After this crude alignment, active alignment is done by optimising the power from the output fibre directly connected to the power meter.

While the fibre coupler has considerable advantages for easy measurements, it has only a limited bandwidth. In practice, the fibre coupler can be used for wavelengths between $1520nm$ and $1580nm$. Outside this range, the measurement errors become rather large.

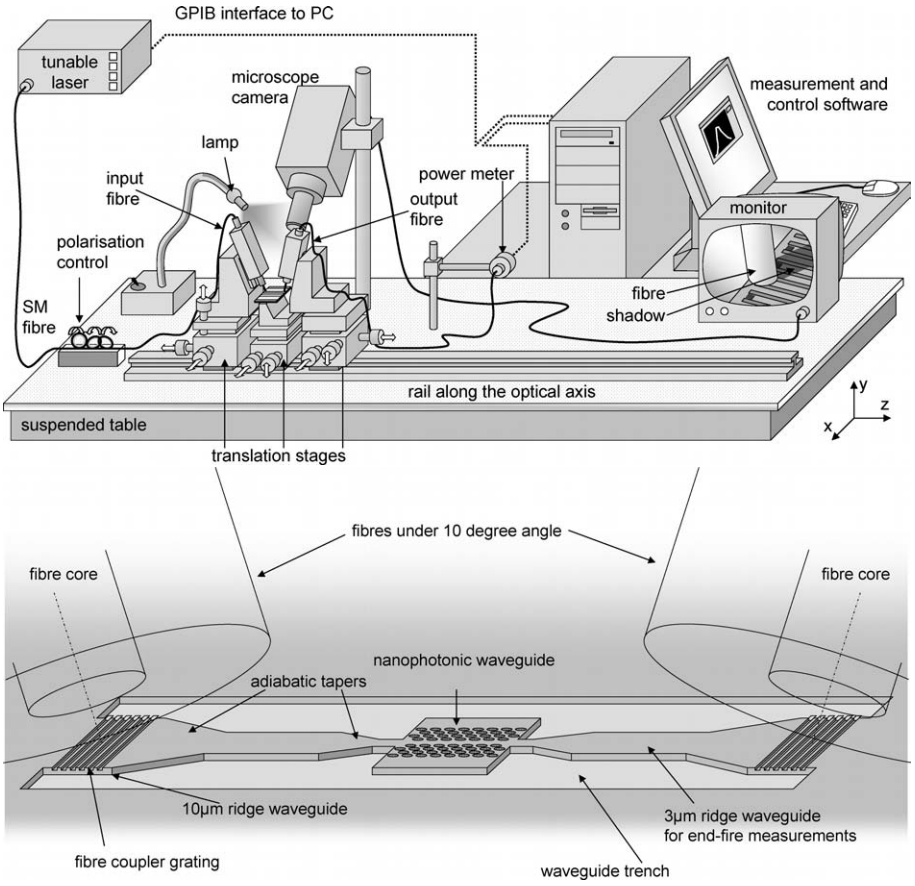

Figure 6.2: The measurement setup with fibre couplers. Light from a tunable laser is coupled into a nanophotonic circuit by means of vertical fibre couplers. The polarisation of the light is controlled by applying torque to the fibre with a set of polarisation control wheels. The fibre couplers couple light into a $10\mu m$ wide ridge waveguide which is tapered down towards a nanophotonic waveguide. The light transmitted through the sample is coupled back to a fibre connected to a power detector. The fibre is positioned visually through a microscope camera by observing both the fibre end and the shadow of the fibre projected by a lamp. The incoupling and outcoupling fibres and sample are put on a translation stage. Both the laser and the power detector are computer controlled through a GPIB interface.

6.2 Measuring Propagation Losses

Both measurements with the end-fire technique and the fibre couplers are transmission measurements: For a given input, the output on the other side of the structure is measured. When the behaviour of all elements in the measurement setup are exactly known, the transmission loss can be extracted directly. However, this is not the case, and there are a number of unknown or fluctuating factors. Both for end-fire and fibre coupler measurements, the incoupling and outcoupling efficiency is not exactly known, and can vary from sample to sample, and sometimes even from waveguide to waveguide. Even more, each structure on a sample contains a number of elements with different losses, and in many cases there are internal interfaces which cause reflections. These reflections create a number of cavities in the structure which result in a very wavelength-dependent transmission, even when the propagation losses of the nanophotonic waveguide are wavelength-independent.

In this section, we will discuss a number of techniques to extract the propagation losses of nanophotonic waveguides from the transmission spectrum of a nanophotonic waveguide. The *cut-back method*, where we subsequently measure the transmission of nanophotonic waveguides with different lengths, forms the basis of all methods. The simplest way to extract the propagation losses is by directly comparing the transmitted power for the different lengths. As we will see further, this method gives a first insight, but has considerable drawbacks.

Instead of directly measuring the power, one can use the properties of *Fabry-Perot cavities*. In the end-fire technique, such a cavity is formed by the cleaved facets. Light propagating in the waveguides is reflected at the facets and circulates back and forth in the waveguide. When the optical roundtrip length is a whole number of wavelengths, all components will interfere constructively and a transmission peak will occur at this wavelength. This effect is stronger when the losses inside the cavity are low. From the contrast of the periodic fringes in this peaked spectrum, we can therefore extract the loss in the whole cavity very accurately [102].

However, when the structure consists of various sections, light can be reflected at the interface. This creates a number of Fabry-Perot cavities, and the transmission spectrum is much more difficult to interpret than that of a single cavity. We will discuss a technique that makes it possible to partially isolate the contributions of the various cavities in

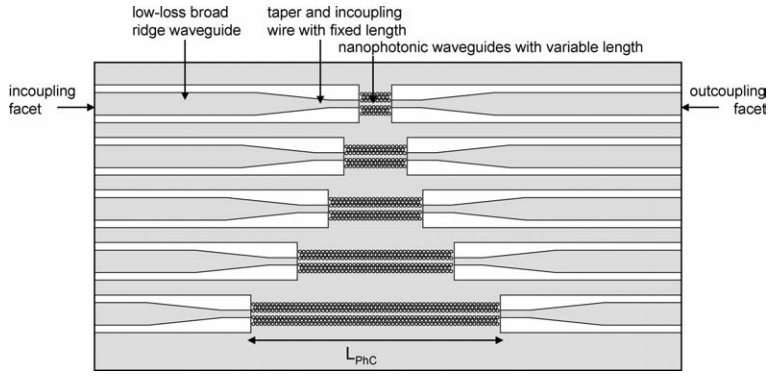

Figure 6.3: Structures for cut-back measurements. The length of the nanophotonic waveguide is increased from one structure to the next at the cost of a broad ridge waveguide with low propagation loss. When all other components are kept the same, the propagation loss of the nanophotonic waveguide can be extracted (when ignoring the propagation loss of the broad waveguide).

the transmission spectrum, and use this information to extract a figure for the propagation losses in the nanophotonic waveguides.

6.2.1 Cut-back Method

To determine the propagation losses of nanophotonic waveguides, we measure the transmission of waveguides with different length. This principle is called the *cut-back* method, because originally (for optical fibres or large planar waveguide) the different waveguide lengths could be made by subsequently cutting a piece of the fibre. In nanophotonic waveguides integrated on a chip, the various lengths are implemented as individual waveguides, as illustrated in figure 6.3.

To extract the propagation losses of the nanophotonic waveguide, we measure the transmission $T = P_{out}/P_{in}$ for each waveguide length L_{phc} and plot the measured values in a $T|L_{phc}$ graph, as in figure 6.4.

6.2.2 Direct Power Measurements

The hardest part of these measurements is to accurately and reproducibly measure the transmission T. The most straightforward way is to measure the power at the detector and divide it through the power emitted by the laser diode. However, the incoupling through a cleaved facet with a lensed fibre is very critical and likely to vary from one

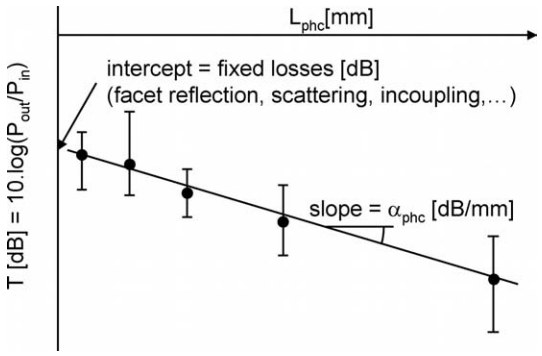

Figure 6.4: Extraction of propagation loss with the cut-back method. When plotting the transmission T in dB ($T\,[dB] = 10\log T$) as a function of L_{phc} length (in mm), the measured points should be on a straight line, of which the slope gives the propagation losses α_{phc} of the nanophotonic waveguide in dB/mm. The intercept with the ordinate axis gives the residual loss in dB.

waveguide to the next. Also, changes in temperature or other environmental factors might influence the overall transmission.

In the structures illustrated in figure 6.3, the light is coupled in and out through a cleaved facet, which has 39% reflection. This forms a Fabry-Perot cavity, where the power oscillates strongly as a function of wavelength, with a ratio between minima and maxima of $3:1$ for the shortest waveguides. In addition, there can be additional reflections at the interfaces between the photonic crystal waveguide and the taper. Trying a fit like in figure 6.4 through a large number of these oscillating signals is bound to introduce an unacceptable error in the extraction of the propagation loss.

Therefore, this measurement technique is only usable with a very lossy Fabry-Perot cavity. This is the case when the nanophotonic waveguides have very high propagation losses, as was the case with our first generation of nanophotonic waveguides. The quality of the Fabry-Perot cavity can also be reduced by eliminating the facet reflections. In layer structures with a low vertical index contrast, this can easily be done with an anti-reflective coating. For SOI, the design of a good AR coating is not as straightforward. Alternatively, the vertical fibre couplers used under an oblique incoupling angle have little reflection, but only a limited bandwidth.

In the case of a broadband light source, direct-power measurement is more widely applicable, as it is possible to average out the Fabry-Perot fringes caused by the facet reflection.

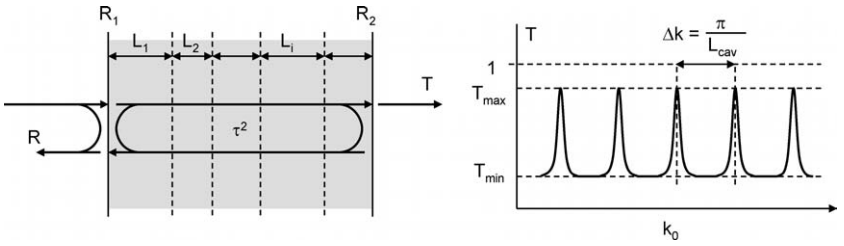

Figure 6.5: A Fabry-Perot cavity and its transmission spectrum. The cavity is formed by a medium (possibly consisting of several sections) and two reflecting interfaces or mirrors at each end. The transmission spectrum from such a cavity has is a periodically peaked function of k_0. The peaks occur when all the circulating components of light in the cavity interfere constructively.

6.2.3 Fabry-Perot Measurements using Fringe Contrast

When we have a single Fabry-Perot cavity, we can use its typical transmission spectrum to determine the losses inside the cavity [102]. Figure 6.5 shows a simple 1-D Fabry-Perot cavity consisting of two interfaces with a reflectivity R_1 and R_2, and a medium with loss τ^2. The medium can be subdivided in various sections i, but we assume no reflections at the internal interfaces. The resulting transmission T of this cavity is

$$T(k_0) = \frac{(1-R_1)(1-R_2)\tau^2}{\left(1-\sqrt{R_1R_2}\tau^2\right)^2 + 4\sqrt{R_1R_2}\tau^2 \sin^2 k_0 L_{cav}}. \tag{6.1}$$

Here, L_{cav} is the *optical length* of the cavity[1]. In a cavity consisting of a number of sections i, L_{cav} would be

$$k_0 L_{cav} = \sum_i k_i L_i, \tag{6.2}$$

with k_i the propagation constant of the light in section i, and L_i the length of the section.

The internal cavity loss τ^2 can be caused by both distributed propagation losses α_i in the various sections and lumped scatterers:

$$\tau^2 = \prod_i e^{-\alpha_i L_i} \prod_j (\chi_j), \tag{6.3}$$

[1] In a uniform medium or a simple waveguide, the *optical length* is equal to the length times the refractive index of the medium (or effective index of the waveguide). However, for structures like tapers, it is hard to define something like an effective index. So, in for the remainder of this text, L_{cav} contains the *average refractive index* of the cavity.

with χ_j the fraction of the light not scattered at each scatterer j.

When we consider the case where the facets have equal reflectivity R the expression for the transmission T is simplified to:

$$T(k) = \frac{(1-R)^2 \tau^2}{(1-R\tau^2)^2 + 4R\tau^2 \sin^2 k_0 L_{cav}}. \tag{6.4}$$

Such a transmission curve is plotted in the right part of figure 6.5. Note that we have plotted the transmission curve as a function of $k_0 = 2\pi/\lambda$ and not of wavelength λ, so we get a periodic transmission characteristic, with a period which is inversely proportional to L_{cav}. Sharper peaks are caused by the lower loss of the cavity due to both internal losses and facet reflections.

From such a periodically peaked spectrum we can now extract the cavity losses. The relevant quantity for this is the *fringe contrast* S:

$$S = \frac{T_{max}}{T_{min}} = \frac{\left(1+R\tau^2\right)^2}{\left(1-R\tau^2\right)^2} \tag{6.5}$$

The value of S can be determined experimentally, and from it we can extract a measure for the total cavity loss $R\tau^2$ due to both internal losses and light escaping at the facets:

$$R\tau^2 = \frac{\sqrt{S}-1}{\sqrt{S}+1}. \tag{6.6}$$

$R\tau^2$ is related to the quality factor Q of the cavity, which is the number of roundtrips in the cavity needed to reduce the intensity of the light with a factor of e:

$$Q = -\frac{2\pi}{\ln\left(R\tau^2\right)}. \tag{6.7}$$

The closer $R\tau^2$ is to 1, the higher the Q of the cavity. $R\tau^2$ can also be expressed in dB:

$$\begin{aligned} R\tau^2\,[dB] &= 10\log\tau^2 + 10\log R \\ &= 10\log\frac{\sqrt{S}-1}{\sqrt{S}+1} \end{aligned} \tag{6.8}$$

Let's consider a Fabry-Perot cavity formed by a broad ridge waveguide tapered down to a photonic wire, as shown in figure 6.6. The cleaved facets function as mirrors, but otherwise there are no internal

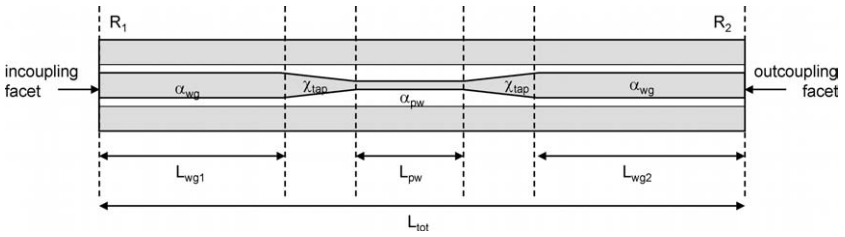

Figure 6.6: A Fabry-Perot cavity with a photonic wire as lossy element. The cavity consists of various sections, each of which will contribute to the cavity losses. The tapers with a fixed width are treated as lumped elements.

reflective elements, as the broad waveguide can easily be tapered adiabatically to a photonic wire. The internal losses in this cavity are formed by the propagation losses of the broad waveguide and the photonic wire. Because the tapers have a fixed physical length[2] independent of the physical lengths L_{wg1}, L_{wg2} and L_{pw} of the different waveguides, we treat them as lumped elements. The internal loss τ^2 now becomes:

$$\tau^2 = e^{-\alpha_{wg}L_{wg}-\alpha_{pw}L_{pw}}\chi_{tap}^2, \tag{6.9}$$

or in dB:

$$\tau^2\,[dB] = -\frac{10}{\ln 10}\left(\alpha_{wg}L_{wg} + \alpha_{pw}L_{pw}\right) + 20\log\left(\chi_{tap}\right). \tag{6.10}$$

When using the cut-back method with these photonic wires, the total physical length of the cavity L_{tot} (not to confuse with the earlier used optical length L_{cav}) is constant:

$$\begin{aligned} L_{tot} &= L_{wg} + L_{pw} + 2L_{tap} \\ &= cst. \end{aligned} \tag{6.11}$$

A cavity with a longer photonic wire will have a shorter section of broad waveguides. Therefore, we can only extract the propagation loss of the wire α_{pw} relative to that of the broad waveguide α_{wg}:

$$\begin{aligned} \tau^2\,[dB] &= -\frac{10}{\ln 10}\left(\alpha_{wg}\left(L_{tot} - 2L_{tap}\right) + \left(\alpha_{pw} - \alpha_{wg}\right)L_{pw}\right) \\ &\quad +20\log\left(\chi_{tap}\right) \qquad (6.12) \\ &= -\frac{10}{\ln 10}\left(\alpha_{pw} - \alpha_{wg}\right)L_{pw} \\ &\quad +20\log\left(\chi_{tap}\right) - \frac{10}{\ln 10}\alpha_{wg}\left(L_{tot} - 2L_{tap}\right). \qquad (6.13) \end{aligned}$$

[2]The optical length of the tapers can be wavelength-dependent.

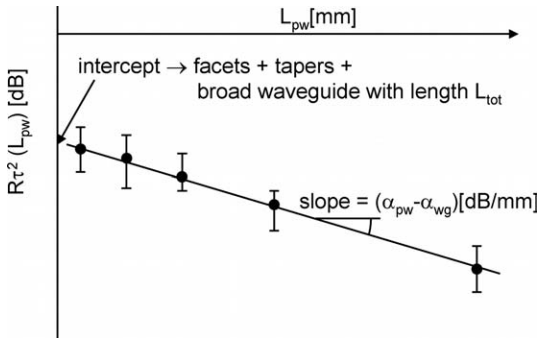

Figure 6.7: Extraction of propagation loss with the cut-back method, but using the $R\tau^2$ extracted from the cavity transmission spectrum instead of the measured power. When plotting this measure for the cavity quality $R\tau^2$ in dB as a function of L_{phc} length (in mm), the measured points should again be on a straight line, of which the slope gives the propagation losses α_{phc} of the nanophotonic waveguide in dB/mm. The intercept with the ordinate axis is given by $10 \log R$ and the residual loss in τ^2 (in dB).

When we now plot the measured values of $R\tau^2\,[dB]$ as a function of $L_{pw}\,[mm]$ (see figure 6.7) we can extract the excess propagation loss $\alpha_{pw} - \alpha_{wg}$ in $[dB/mm]$ from the slope of the linear fit through the measurement points. Typically, α_{wg} is much smaller than α_{pw}, so the slope is a very good measure for the propagation losses of the photonic wire. The intercept of the line with the ordinate axis contains the other internal cavity losses and the facet reflection.

Note that measuring the fringe contrast S requires a fine wavelength scan, in order to establish the exact value of T_{max}, especially for cavities having low loss. Nevertheless, by measuring the fringe contrast, we can extract the cavity losses independently of the incoupling and outcoupling efficiency of the measurement setup, significantly boosting the reliability of the measurements.

While this method allows us to measure the propagation losses of photonic wires very accurately, it has its limitations. Firstly, the photonic wire should be single-mode. If not, the interference between the two or more modes give rise to an additional modulation of the spectrum. For broader wires, which support multiple modes, this problem can be solved by including a short narrow section to filter out the higher-order modes.

Also, this method only works with a single cavity. When the structure contains reflecting interfaces inside the cavity, additional cavities

are formed, each of which gives rise to a Fabry-Perot spectrum. These overlaid spectra cannot be interpreted as straightforward as the one of a single cavity. Because in photonic crystal waveguides the interface between the crystal and the wire always has some reflection, we cannot use this technique. In the next sections, we will discuss another method which will allow us to separate the contributions of the various cavities.

For the structures fabricated with deep UV lithography, the facets are typically $4-6mm$ apart. Typically, the effective index of the broad incoupling and outcoupling waveguides is around 2.8. This creates an optical cavity length of $10mm$ or more. To measure the fringes of such a cavity with sufficient accuracy, a wavelength scan with a step of at most $0.05nm$ is needed. Scanning the whole measurement range from $1500nm$ to $1640nm$ with this resolution requires several hours, during which the setup, including incoupling and outcoupling, should be kept stable.

6.2.4 Fourier Analysis of Fabry-Perot Cavities

Coupling light into a photonic crystal waveguide is not as straightforward as coupling it into a photonic wire. Because there is always some reflection at the interface between the photonic crystal and the incoupling wire, multiple cavities are formed between the incoupling and outcoupling facets. Figure 6.8 shows the simple case of a straight photonic crystal waveguide, where there are already six first-order cavities. More cavities are formed when combining two or more of these six cavities [103].

Before analysing such a multi-cavity system, we will look again at the single cavity from figure 6.5. When $R_1 = R_2 = R$, the transmission as a function of k_0 becomes

$$T(k_0) = \frac{(1-R)^2 \tau^2}{(1-R\tau^2)^2 + 4R\tau^2 \sin^2 k_0 L_{cav}}. \tag{6.14}$$

This function of k has a period of π/L_{cav}, or a frequency of L_{cav}/π. Therefore, we can use the optical length L as the measure for 'frequency' in the Fourier transform:

$$F\{T\}(L) = \int T(k)\, e^{j\pi L k} dk \tag{6.15}$$

If we consider $T(k)$ to be an infinitely periodic function, the Fourier transform is a series of delta peaks at multiples of L_{cav}, convolved with

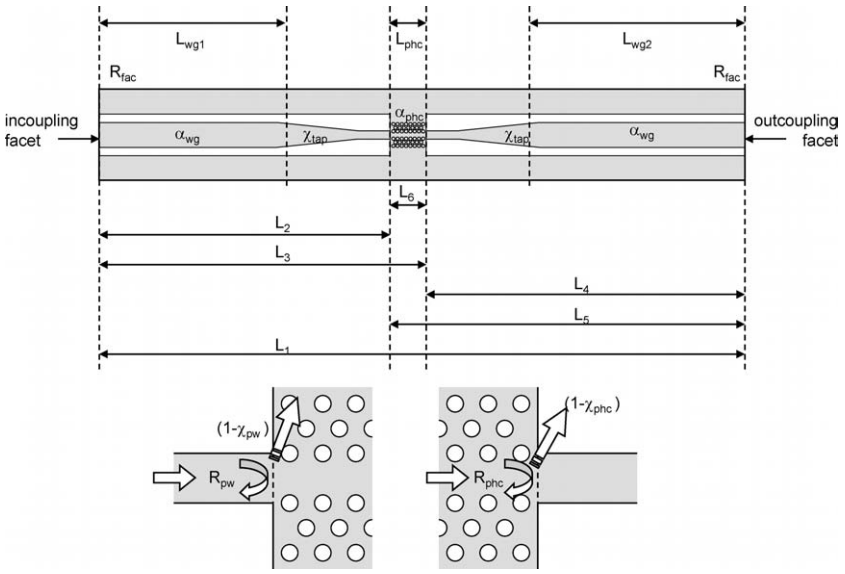

Figure 6.8: A multiple Fabry-Perot cavity system formed by a photonic crystal waveguide. The additional cavities are formed by the reflection at the interface between the photonic crystal and the photonic wire. In the first order, 6 cavities are formed, with lengths L_1 through L_6. To complicate the matter even more, the reflection at this interface is not necessarily the same in both directions.

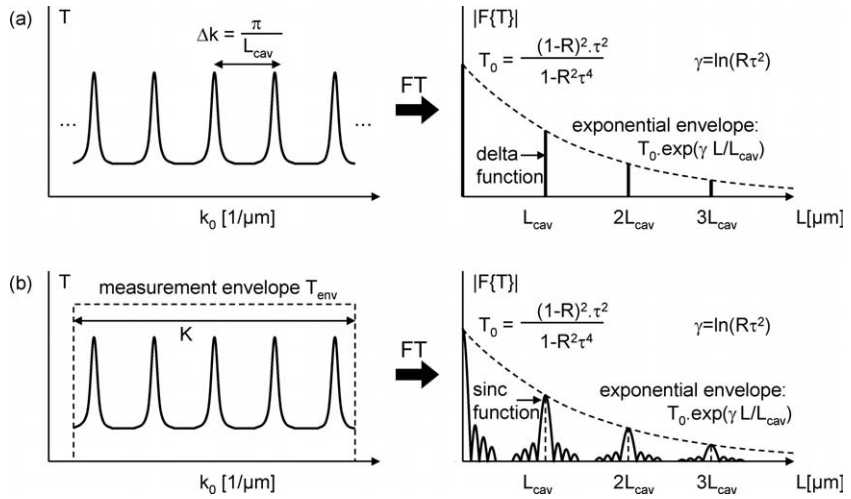

Figure 6.9: The Fourier transform of a single Fabry-Perot Cavity. The periodically peaked transmission spectrum is transformed into a series of peaks with an exponentially declining envelope. (a) The infinite transmission spectrum gives rise to a series of delta functions. (b) For a finite measurement window, the delta functions are smeared out as sinc functions.

an exponential envelope function. This is illustrated in figure 6.9a. When we calculate the Fourier transform analytically, discard the complex phase and look only at the absolute value, we get the following expression for $|F\{T\}|$ [104]:

$$|F\{T\}(L)| = \sum_i \delta(L - iL_{cav}) * \left(T_0 e^{\gamma \frac{L}{L_{cav}}}\right), \tag{6.16}$$

with

$$T_0 = \frac{(1-R)^2 \tau^2}{1 - R^2\tau^4}, \tag{6.17}$$

the average power over the spectrum and

$$\gamma = \ln\left(R\tau^2\right) \tag{6.18}$$

describing the exponential decay in each subsequent peak.

Loss Extraction from a Single Peak

From the height of the peaks in the Fourier transform we can directly extract the losses inside the cavity. The central peak at $L = 0$ is the most

prominent, and has a peak height P_0 of

$$P_0 = \frac{(1-R)^2 \tau^2}{1 - R^2\tau^4}. \tag{6.19}$$

From this expression we cannot extract $R\tau^2$ analytically as we could with the fringe contrast method. However, when we know R in advance, we can extract the value for τ^2 from the peak height:

$$\tau^2 = -\frac{(1-R)^2 + \sqrt{(1-R)^4 + 4R^2P_0}}{2R^2P_0}. \tag{6.20}$$

Plotting the measurement of $\tau^2\,[dB]$ as a function of waveguide length will then directly yield the propagation losses. However, the facet reflection R should be determined in advance either in another experiment or by simulation.

As we will see further, this technique can be used in the situation of multiple cavities as well. However, in a multi-cavity system the peak at $L = 0$ is a contribution of various cavities, and its height not straightforward to interpret. The first peak, at $L = L_{cav}$, however, can in most cases easily be isolated. We can also use the height of this peak to extract the internal cavity loss τ^2. The height of the first peak P_1 is given by

$$P_1 = R\tau^2 \frac{(1-R)^2 \tau^2}{1 - R^2\tau^4}. \tag{6.21}$$

When we solve this to τ^2, we get

$$\tau^2 = \sqrt{\frac{P_1}{R\,(1-R)^2 + R^2P_1}}. \tag{6.22}$$

This expression for τ^2 is even more simple than the one extracted from the peak at $L = 0$.

For cut-back measurements, this extraction of τ^2 is a better alternative to the direct power measurements. Its main advantage is that it filters out the oscillating Fabry-Perot characteristic.

Unlike the fringe contrast technique, this method is still sensitive to variations in incoupling and outcoupling efficiency between waveguides, as it is still normalised to the absolute power level. However, it is preferable to the approach of direct power measurements.

Loss Extraction from the Ratio Between Peaks

Instead of using a single peak, we can measure the ratio between two subsequent peaks to extract the losses. For a single Fabry-Perot cavity, the ratio between two adjacent delta peaks in the Fourier transform is $R\tau^2$, and therefore a direct measure for the cavity quality. This can be intuitively interpreted as follows. The peak at $L = 0$ represents the light passing directly through the cavity. Each subsequent peak represents light that has made an additional roundtrip in the cavity. For a high-quality cavity, the roundtrip loss is not large, and therefore the peaks will only slowly diminish. This translates in a value of $R\tau^2$ close to 1. For a bad cavity, the peaks will rapidly dwindle to zero, as most light is lost in the first few roundtrips through the cavity.

As with the fringe-contrast method, we can use the values of $R\tau^2$ for different waveguide lengths to determine the propagation losses of a photonic wire, irrespective of the incoupling efficiency.

A Finite Measurement Window

However, when we are working with a measured spectrum of finite width, the response is no longer a series of delta function, but the individual peaks will have a finite width. In figure 6.9b, the finite measurement window places a square envelope over the transmission curve. The resulting Fourier transform is now a convolution of a series of sinc functions with the same exponential envelope function:

$$|F\{T\}(L)| = \sum_i \frac{\sin\left(K\left(L - iL_{cav}\right)\right)}{K\left(L - iL_{cav}\right)} * \left(T_0 e^{\gamma \frac{L}{L_{cav}}}\right). \tag{6.23}$$

The width of the individual peaks is inversely proportional to the width of the measurement window K. Also, the sinc function has considerable sidelobes. For short cavities or narrow measurement windows, these sidelobes might interfere with neighbouring peaks, distorting the spectrum.

Because of its large sidelobes, a sinc function is ill suited as a measurement window. Instead, one can modulate the measured signal with an artificial window function. A number of these window functions are shown in figure 6.10. The surface area under each window is the same, which results in peaks of the same height. While in theory, a pure gaussian window would give the narrowest peak without sidelobes, this would also require an infinite measurement range. Figure 6.10b and

6.10c show two truncated gaussian functions with different width, and both still have significant sidelobes. With a finite measurement range, a *hanning* window gives the best reduction of the sidelobes, but with a broader central peak.

The width of the central peak determines the smallest differences in cavity length that can be resolved. Typically, we can resolve two different peaks in the Fourier spectrum when the center of the second peak is above the first minimum of the first peak. The width of the peaks is shown in figure 6.10 expressed in π/K. For a measurement range of $1.5\mu m$ to $1.6\mu m$, π/K is $12\mu m$. Using a *hanning* window, this means that two peaks should at least be $25\mu m$ apart, which gives us the shortest cavity length we can resolve.

However, using such a broad wavelength range is not always advisable. As nanophotonic waveguides seldom have a linear dispersion, the optical cavity length will change as a function of wavelength. When transforming over a broad window, the peaks in the Fourier spectrum will be smeared out. Therefore, with dispersive waveguides, it is better to keep the window as narrow as possible. This, of course, limits the shortest cavity length that can be resolved.

Note that to get well defined peaks in the Fourier transform, a very fine scan of L is needed. This excludes the use of the Discrete Fourier Transform and its efficient Fast Fourier Transform (FFT) implementation. Instead, a numerical integration of the continuous Fourier transform should be used.

Multiple Fabry-Perot Cavities

Because the transmission spectrum of such a multiple-cavity system is very complex and contains both short-range as long-range oscillations, it is hard to interpret directly. However, we have seen that the periodic spectrum of a Fabry-Perot cavity has a period which is inversely proportional to the optical cavity length. A multi-cavity transmission spectrum will therefore be a superposition of periodic functions with a different period for each cavity. These periodic components will translate to different peaks in the Fourier transform [104, 105, 103]. Figure 6.11 shows that the Fourier transform of the infinite transmission spectrum contains a number of peaks, representing the various cavities. Because of the strong reflection at the incoupling and outcoupling facets, the highest peak is at $L = L_1$. Depending on the reflection at the interfaces between photonic crystal and photonic wire, the other first-order

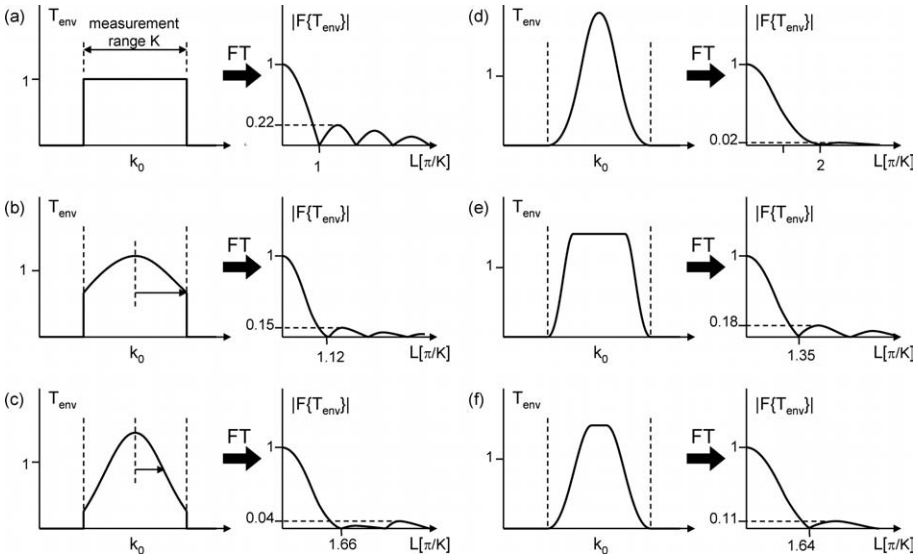

Figure 6.10: Window functions for the Fourier transform. When using a measurement window of a finite width, the peaks in the Fourier transform will no longer be delta functions, but become smeared out. (a) For a rectangular measurement window, the peaks will be sinc functions with strong sidelobes. (b,c) When using an infinite gaussian, the peaks will no longer have sidelobes. However, this is not possible as we only have measurements over a finite wavelength range. Truncated gaussian windows do have sidelobes, but less severe than those of the rectangular window. (d) A *hanning* window, consisting of one period of a cosine function, is the best option for suppressing the sidelobes, but the peak is twice as broad as for the rectangular window. (e,f) *Tukey* windows are rectangular windows, but at the edges they go to zero in the same way as a hanning window. By playing with the width of the tukey window, the peak width and the sidelobe height can be adjusted. In order to keep the peak height constant, the area under the window function should be kept constant.

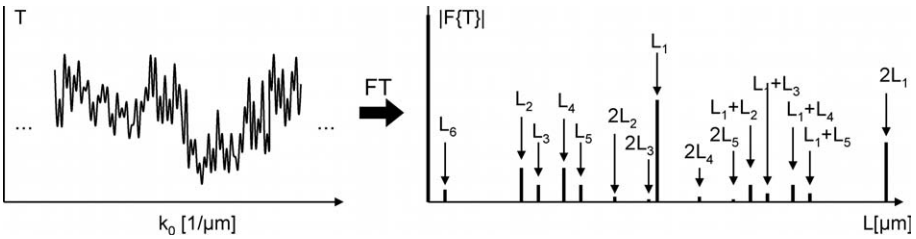

Figure 6.11: The Fourier transform of multiple single Fabry Perot Cavity. The interference of the multiple cavities of figure 6.8 give rise to an irregular spectrum with both long-range and short-range modulation. After Fourier transformation, the individual cavities appear as peaks in the transformed spectrum.

cavities appear in various strengths and even higher-order cavities may appear.

Just as in the single-cavity case, the delta peaks will become finite peaks when a finite measurement window is used. As already shown above, the width and shape of the measurement window determines how well the individual peaks can be resolved. In practice, the most critical situation occurs when measuring the transmission of short photonic crystal waveguides. This makes the difference between L_2 and L_3, respectively L_4 and L_5 very small, effectively joining the two peaks in the transformed spectrum.

Just as with a single cavity, we can extract the internal cavity losses τ^2 from the single peak of cavity L_1. Again, we need to make a estimate of the facet reflectivity. In τ^2, not only the propagation losses of the various waveguides are included, but also the additional reflections and scattering at the interface between the photonic crystal and the photonic wire:

$$\tau^2 = e^{-\alpha_{wg}L_{wg}-\alpha_{phc}L_{phc}}\chi_{tap}^2\left(1-R_{pw}\right)\chi_{pw}\left(1-R_{phc}\right)\chi_{phc}. \qquad (6.24)$$

Note that the reflection and scattering at the interface between the photonic crystal and the photonic wire are not necessarily identical in both directions. However, due to reciprocity, they are related as

$$\left(1-R_{phc}\right)\chi_{phc} = \left(1-R_{pw}\right)\chi_{pw}. \qquad (6.25)$$

If R_{pw} and R_{phc} are much smaller than R_{fac}, the first peak of cavity L_1 will be much higher than the surrounding peaks. Therefore, the overlap of neighbouring peaks has little effect on the height of this

peak. This means we can afford to use broad peaks, and therefore a narrow measurement window. This is useful, because in photonic crystal waveguides the propagation loss can be very wavelength-dependent. By scanning the transmission spectrum with a narrow window, we can calculate τ^2 as a function of wavelength.

Figure 6.12 shows an example of this technique. Even though the transmission spectrum has a significant dip, this dip does not show up in the extracted cavity loss τ^2. This is because the dip in the transmission curve is caused by the slow Fabry-Perot oscillation of cavity L_6 and not because of a lower transmission of the photonic crystal waveguide. This immediately shows the advantage of this technique over measuring the power directly. However, it is still subject to variations caused by irreproducible incoupling and outcoupling.

Instead of using a single peak, we can also use the ratio of the first two peaks of cavity L_1. To do this, however, the losses in the cavity should not be too large, as we should be able to resolve the second peak. This method gives us the cavity losses while eliminating the uncertainty of the incoupling efficiency. Therefore, where possible, this technique should be used to analyse the measurements

Sample Preparation

While analysing the peaks in the Fourier spectrum is the best method of analysing measurement results, some caution is needed before applying it, especially in a multi-cavity system. When preparing a sample with cleaved facets, care should be taken that the cleaves are not symmetric with respect to the photonic crystal waveguide. In other words, $L_2 \neq L_4$. In the case of short photonic crystals, the second order peaks of these cavities, and of the cavities L_3 and L_5 will closely overlap the peak of cavity L_1, making it impossible to extract cavity losses of L_1.

6.2.5 Conclusion

We have discussed a number of ways to analyse the measurement spectrum to extract the propagation losses of nanophotonic waveguides. In all cases, we have used the cut-back method, analysing the losses for different waveguide lengths. While the most straightforward way to extract the propagation losses is to plot directly the measured power as a function of waveguide length, this introduces large errors, mainly due to Fabry-Perot oscillations and irreproducible incoupling.

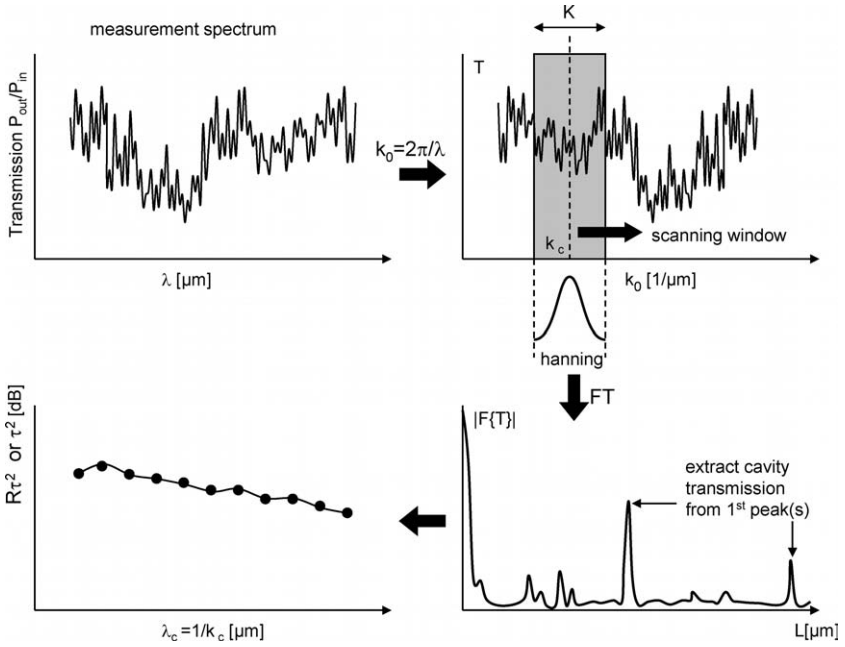

Figure 6.12: Extracting wavelength-dependent propagation losses with a local Fourier transform. The measurement spectrum (artist's impression) as a function of λ is converted to a spectrum as a function of $k_0 = 2\pi/\lambda$. A hanning window is placed on the spectrum before is transformed. As the hanning window can be much narrower than the total measurement spectrum, it can be scanned over the entire measurement range. From these local Fourier transforms, either τ^2 or $R\tau^2$ can be extracted respectively from directly measuring one peak, or comparing two peaks. This results in a cavity loss plot as a function of wavelength. In the example there is a dip in the measurement spectrum which does not occur in the extracted cavity loss. This is because the dip is the result of a slow oscillation from the cavity formed by a short photonic crystal waveguide.

Analysing the fringe contrast in spectrum from the Fabry-Perot cavity formed by the cleaved facets gives a much more accurate estimate of the losses, irrespective of the incoupling efficiency However, it can only be used in a clean cavity, without internal reflections.

When there are internal reflections, like in the case of a photonic crystal waveguide, we can analyse the transmission spectrum in the Fourier domain. This translates into a number of distinct peaks corresponding to the various cavities in the system. By measuring the peak height directly we can extract the power, eliminating the uncertainties due to Fabry-Perot oscillations. However, it does not solve the problem of irreproducible incoupling. Still, this technique is superior to direct power measurements

By taking the ratio between two peaks in the transformed spectrum, we can extract the cavity loss of the Fabry-Perot cavity, just as in the fringe contrast method. This method has the strengths of the fringe contrast method, with the advantage that it can be used in a multi-cavity system. The only requirement is that the peaks can be sufficiently resolved, and that the peaks of the subcavities do not strongly interfere with the peak of the main cavity.

6.3 Fabricated Components

In the course of this work, we have fabricated a large number of nanophotonic components. Because deep UV lithography is a mask-based fabrication process, these designs were grouped in time on four masks. Appendix A gives an overview of the four masks. These masks include both process validation structures and functional components, like photonic crystal waveguides, photonic wires and fibre couplers.

6.3.1 Device Generations

Once the basic process characteristics were determined, functional devices could be fabricated. Each generation of devices improved on the previous generations by using different substrates or etch processes. Because $248nm$ lithography proved to be adequate for this purpose, we used the same lithography process for all generations. In parallel with the process development, designs were also improved and included on the new masks. Table 6.1 gives an overview of the device generations fabricated with deep UV lithography for high-contrast waveguides. Table 6.2 gives the various generations of vertical fibre couplers.

Generation		SOI		Etch		Masks				Remarks
		Si	Ox	Si	Ox	1	2	3	4	
1	mar/02	205	400	205	400	X				Thin ox.
2	feb/03	220	1000	220	800	X	X			Roughness
3	jun/03	220	1000	220	none		X			
4	aug/03	220	1000	220	none			X		2-step
5	early 04	220	1000	220	none				X	2-step

Table 6.1: Overview of the generations of high-contrast SOI waveguides, like photonic wires and photonic crystal waveguides, with the basic elements of the fabrication process. Layer thickness and etch depth are in nm.

Generation		SOI		FC etch		WG etch		Masks				Remarks
		Si	Ox	Si	Ox	Si	Ox	1	2	3	4	
1	mar/02	205	400	50	none	50	none	X				1-D, thin ox.
2	feb/03	220	1000	50	none	50	none	X				1-D
3	aug/03	220	1000	50	none	50	none		X			2-D
4	aug/03	220	1000	50	none	220	none			X		1-D, 2-step
5	early 04	220	1000	50	none	220	none				X	2-D, 2-step

FC etch: etch depth of the fibre coupler
WG etch: etch depth of the waveguides

Table 6.2: Overview of the generations of vertical fibre couplers in SOI fabricated with deep UV lithography. In the later generations, fibre couplers are fabricated together with high-contrast waveguides in a 2-step process. Layer thickness and etch depth are in nm.

6.4 Generation 1: SOI with Thin Oxide

Following satisfactory lithography and etch tests, a first batch of structures, implemented on PICCO_1 mask, were fabricated in SOI with a top Silicon layer of $205nm$ and a buried oxide layer of $400nm$. As already discussed in chapter 4, this oxide is too thin to provide adequate isolation from the substrate. This translates in high propagation losses due to substrate leakage.

The structures from the PICCO_1 masks were printed with a variable exposure dose, resulting in a large parameter space of lattice pitch, hole diameter and line width. After lithography, the structures were etched with both a Silicon and an oxide etch. Because the oxide was thin, we could easily etch completely through the oxide. Some fabricated structures are shown in figure 6.13. We can see that the fabrication quality is very good.

First measurements quickly showed that the high substrate leakage completely dominates the losses in such a way that the quality of the Fabry-Perot cavity is extremely low. Therefore, direct-power measurements provide the only reliable way of measuring.

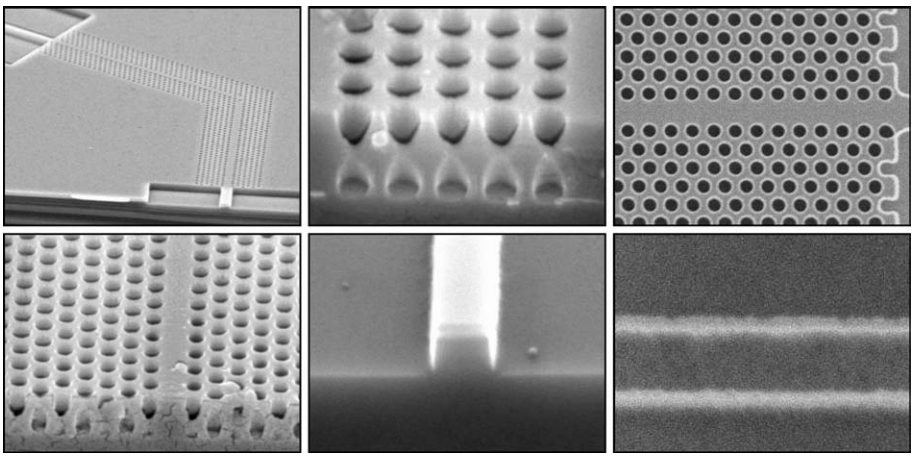

Figure 6.13: Nanophotonic structures in SOI with a thin oxide buffer. Top row: a 60° bend in a photonic crystal, photonic crystal holes on an obliquely cleaved facet, a W1 photonic crystal waveguide. Bottom row: A cleaved W1 waveguide, a cleaved photonic wire, a photonic wire in top-down.

6.4.1 Determining Substrate Leakage

Before measuring nanophotonic waveguides, we tried to measure the substrate leakage of a slab waveguide and compare it to the theoretical calculations of chapter 4. For this, we used a $4\mu m$ wide ridge waveguide, which behaves much like a slab waveguide. Figure 6.14 gives the transmission P_{out}/P_{in} in dB for such a waveguide. As we expect, the waveguide losses increase with the wavelength, as can also be seen in figure 4.23. According to the simulations, the substrate leakage for a slab should rise from $7dB/mm$ at $1520nm$ to $13.5dB/mm$ at $1620nm$. If we divide the measured loss in dB by length of the sample ($2mm$), we do find that the difference between the minimum and the maximum is $6.5dB$, exactly as predicted by the theory. If we now assume the theory is correct, we can calculate the residual losses caused by incoupling and outcoupling. We find that these are overall between $15.5dB$ and $17dB$, with a minimum incoupling loss around $1550nm$, the wavelength at which we optimised the incoupling. The contributions due to substrate leakage and other losses are also indicated in figure 6.14.

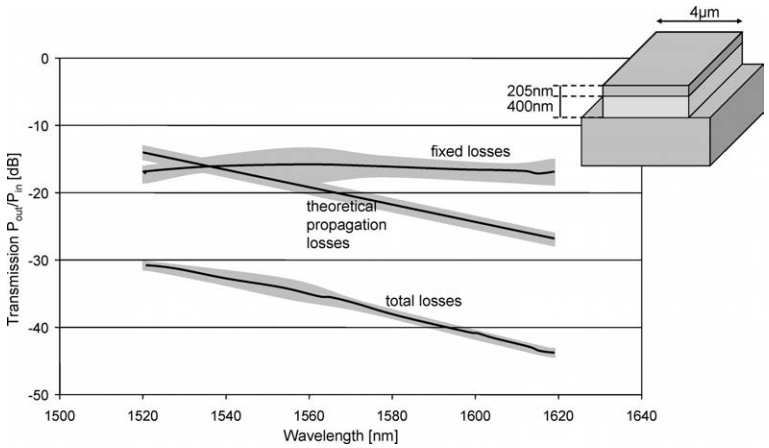

Figure 6.14: Transmission of a $4\mu m$ broad ridge waveguide with thin oxide. The total loss is comprised of propagation losses and fixed losses due to incoupling and outcoupling. The wavelength dependency of the losses matches very well the theoretical model of substrate leakage.

6.4.2 Photonic Wires

We measured the propagation losses of photonic wires with the cut-back method. For this, we had structures with a wire length of $0\mu m$ (only the tapers), $100\mu m$, $200\mu m$ and $400\mu m$. By varying the exposure dose, we could also make wires of various width.

Note that with the cut-back method, we measure the relative propagations loss $\alpha_{pw} - \alpha_{wg}$ of the wire with respect to the broad ridge waveguide. Here, the losses of the ridge waveguide are not negligible, so we have to add them to the the measured values.

Figure 6.15 shows the propagation loss α_{pw} of photonic wires of various width. We can see that, as expected, the propagation losses increase with wavelength more strongly than for the ridge waveguide. Also, as the wire gets narrower, the propagation losses become very large. Not plotted are the results of a photonic wire with a width of $500nm$, which has a propagation loss of $130dB/mm$ at $1550nm$ wavelength.

6.4.3 Photonic Crystal Waveguides

The photonic crystal waveguides from the first generation of components were based on designs readily available in literature [7, 52]. However, because of the high substrate losses, it is not straightforward to

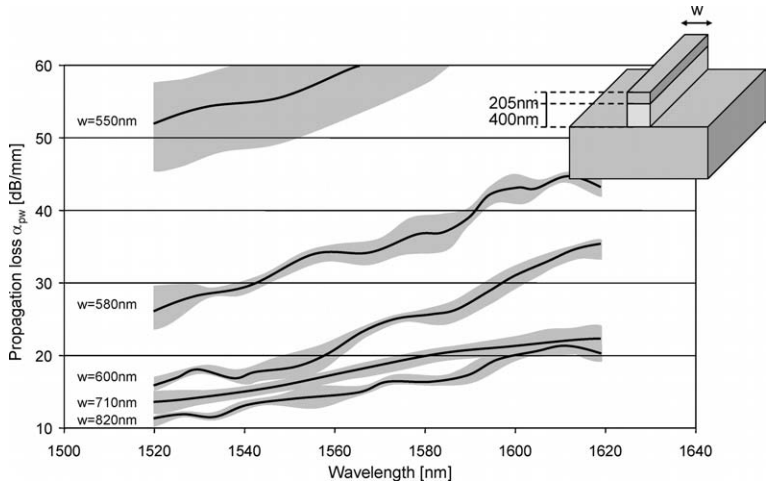

Figure 6.15: Propagation loss of photonic wires with thin oxide. As with the broad ridges, the losses of the photonic wires increase with the wavelength. Also, the losses increase when the wire gets narrower.

interpret the transmission spectra. Because the multiple-cavity system created by the photonic crystal waveguide has little influence due to the high losses, we have to use direct-power measurements.

Figure 6.16 shows the propagation loss of a W1 photonic crystal waveguide with a lattice pitch $a = 460nm$ and a hole diameter $\emptyset = 250nm$. When compared to the band diagram, we can discern three regimes in the loss curve. For short wavelengths, the guided mode is above the light line, with losses of up to $140dB/mm$. In the region between $1550nm$ and $1615nm$, the guided mode is below the light line, and the propagation losses drop down to $100dB/mm$. In this region we can also see that the propagation losses increase with the wavelength due to the substrate leakage, just as with photonic wires. At the side of the loss curve the losses increase rapidly at the location of the mini-stop band created by the interaction of the TE mode and the TM mode.

6.4.4 Generation 1: Conclusion

For our first fabricated devices, the wafer choice determined the performance of our waveguides. With the thin oxide, the substrate leakage is by far the dominating loss mechanism. For broad ridge waveguide, the measured propagation losses are in very good agreement with the theoretical calculations. For single mode photonic wires ($w_{pw} < 560nm$) we

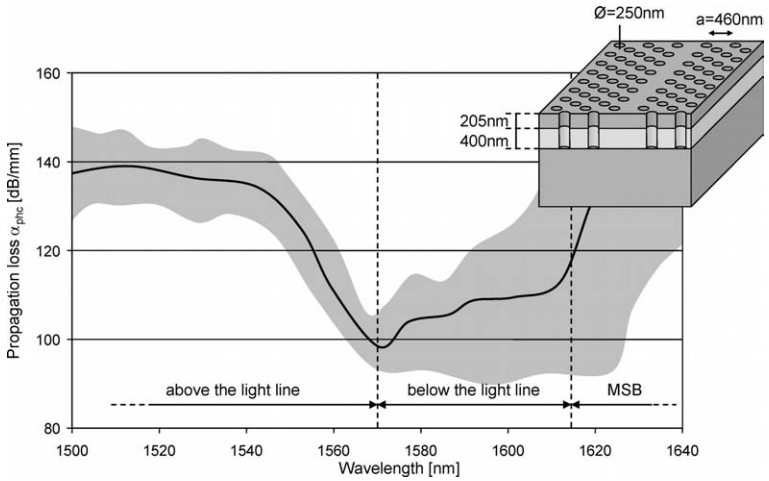

Figure 6.16: Propagation loss of a W1 photonic crystal waveguide with thin oxide. There are three regimes indicated, from left to right: The guided mode above the light line, the guided mode below the light line, and a mini-stopband (MSB), where there is no guided mode. The large errors on the right side are caused by the very low absolute transmission.

have measured propagation losses of $52dB/mm$ at $1520nm$ and increasing towards longer wavelengths. For photonic crystal waveguides, the propagation losses of a W1 waveguide were at least $100dB/mm$.

6.5 Generation 2: Deeply etched SOI

In order to eliminate the losses caused by substrate leakage, wafers with thicker oxide were purchased. In the mean time, a new mask was designed with better designs, including long nanophotonic waveguides for accurate loss measurements.

As discussed in chapter 5, we also developed an etch process to etch the oxide deeper than with the first set of wafers. However, the deep etching caused an unacceptable amount of roughness on the sidewalls of the structures, as can be seen in figure 6.17. As we have seen in chapter 4, this sidewall roughness is a major cause of scattering losses in nanophotonic waveguides.

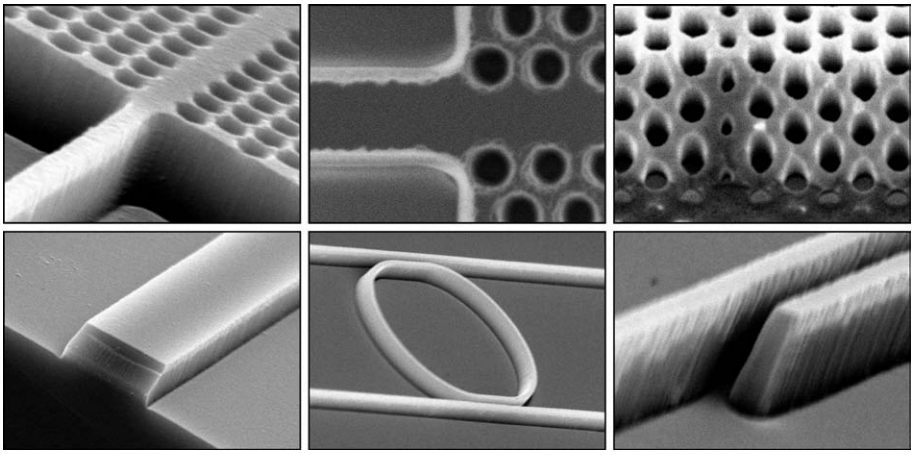

Figure 6.17: Nanophotonic structures in SOI with a thick oxide buffer and a deep etch. Top row: A wire with a W1 photonic crystal waveguide, the same in top-down view, an obliquely cleaved W1 waveguide with small center holes. Bottom row: A cleaved facet of a $3\mu m$ wide ridge, the beginning of a directional coupler, a racetrack resonator. Note the visible roughness on the sidewalls of these structures.

6.5.1 Photonic Wires

Because the loss mechanism of substrate leakage is negligible in these structures, the losses should be low enough to allow Fabry-Perot measurements either based on fringe contrast or with the Fourier transform. For photonic wires, we expect a clean periodic transmission spectrum of a single Fabry-Perot cavity. However, the transmission spectra, like the one of a $500nm$ wide wire in figure 6.18, looks more like the spectrum of multiple cavity system. This is confirmed by the Fourier transform, also shown in figure 6.18. The length of the multiple cavities suggest a reflective element at the location of the photonic wire, indicated by L_{def}. However, inspection with a microscope didn't show any defect in the waveguide.

Luckily, we can still determine the propagation losses of the photonic wires. From the Fourier transforms in figure 6.18 we can extract the propagation losses from the peak at $L = L_{cav}$. The results are plotted in figure 6.19. They indicate a propagation loss of approximately $34dB/mm$ for a $500nm$ wire, and $13dB/mm$ for a $550nm$ wire. We did similar measurements for wires of different width. As expected, the wires have more loss when they are narrower, which agrees with the idea that the majority of the losses is now caused by sidewall rough-

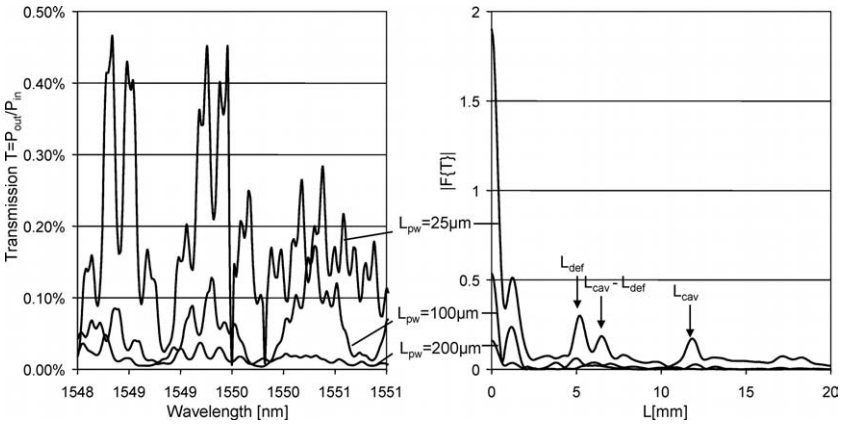

Figure 6.18: Transmission spectra and Fourier transform of deeply etched photonic wires with a width of $500nm$ and various lengths. For clarity, only the wire lengths $L_{pw} = 25, 100, 200\mu m$ are plotted. Although no reflection inside the cavity between the facets is expected, we can see a distinct peak at $L - L_{def}$ in the Fourier transform of the spectra.

ness. For a $600nm$ wire, which is already multi-mode, we find propagation losses of around $6dB/mm$.

6.5.2 Photonic Crystal Waveguides

Photonic crystal waveguides are more difficult to characterise. There is the additional reflection at the interface with the wire, they are more dispersive and sometimes multi-mode. As we have seen in chapter 3, many waveguide designs also have a first-order guided mode in the PBG. To excite these antisymmetric modes with our symmetric incoupling spot from the lensed fibre, we have designed a mode converter. Figure 6.20 shows how an 11° abrupt bend in the broad ridge waveguide mixes the even ground mode into a $50\%/50\%$ combination of the ground mode and the first order mode. Of course, the incoupling wire should also be somewhat broader to support both modes.

On PICCO_2, we have implemented a large number of the photonic crystal waveguide designs discussed in chapter 3. All structures were put on the mask with conventional incoupling structures to excite the ground mode, and with the 50% mode mixer. Comparing the transmission spectrum of both structures can give some insight on the relative propagation loss of the even and the odd modes.

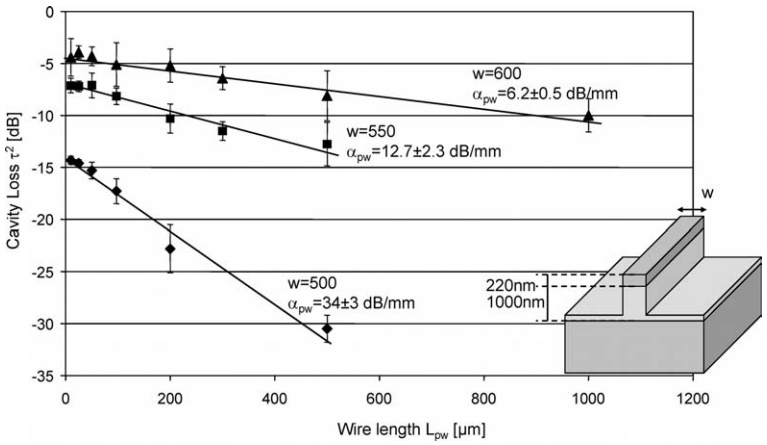

Figure 6.19: Propagation losses of deeply etched photonic wires with various widths. The losses are extracted from the slope of the straight line fitted through the peak height of the peak at $L = L_{cav}$ in the Fourier transform of the transmission spectrum (see figure 6.18). The propagation losses increase for narrower wires.

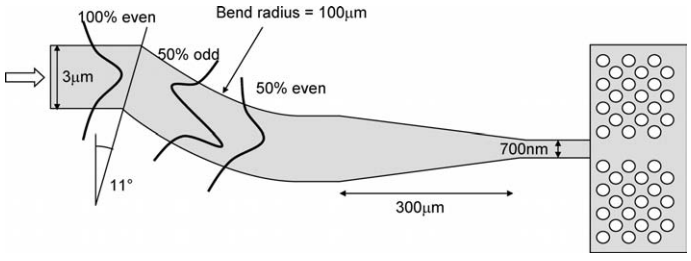

Figure 6.20: Mode mixer to excite asymmetric photonic crystal modes. At the 11° tilt the ground mode is converted into a $50\%/50\%$ combination of the ground mode and the asymmetric first-order mode.

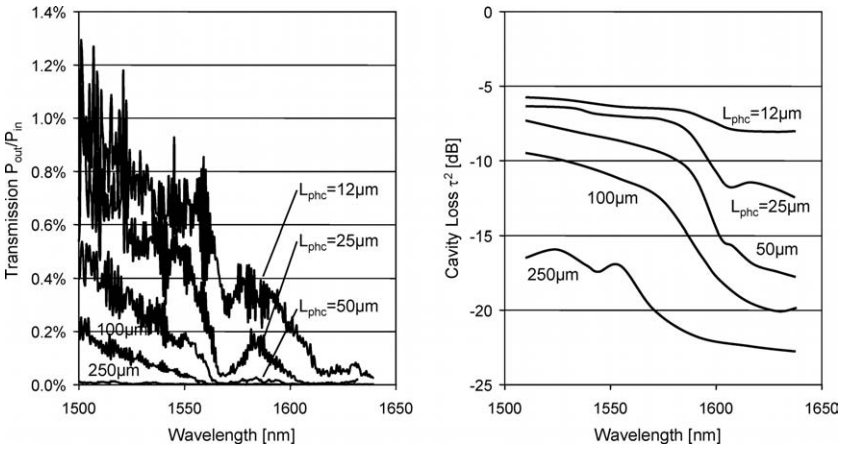

Figure 6.21: Transmission spectra of deeply etched W1 photonic crystal waveguide with a pitch of $500nm$ and a hole diameter of $337nm$. We used incoupling with the even mode. On these spectra, we have used a local Fourier transform and measured the peak height for $L = L_{cav}$, as illustrated in figure 6.12. From this, we can extract the internal cavity loss τ^2, plotted in the right graph.

Figure 6.21 shows the transmission spectra of W1 photonic crystal waveguides of various lengths. The photonic crystal lattice has a period of $500nm$ and the holes are $337nm$ in diameter. In this example we didn't use the mode mixer, and we coupled in with the ground mode of a $500nm$ wide photonic wire. We can see the transmission drop towards longer wavelengths. We analysed these spectra using the height of the first peak in the Fourier transform of the spectra, and this for a local window which is scanned over the whole wavelength range. The internal cavity loss τ^2 extracted from the peak height is plotted in the right part of figure 6.21.

From the peak heights we can now extract the propagation losses of the photonic crystal waveguides, as well as the cavity losses independent of the waveguide length. We can see that the losses drop sharply near $1570nm$, which matches the mini-stop-band between the PBG-guided TE mode and the index-guided TM mode (see band diagram 3.8 on page 60). For the guided mode, the propagation losses are as low as $21dB/mm$ around $1515nm$.

As another example, we measured a W1 photonic crystal waveguide with small defect holes. In figure 6.23 we see the transmission spectra of four such waveguides with a lattice constant of $500nm$, a lat-

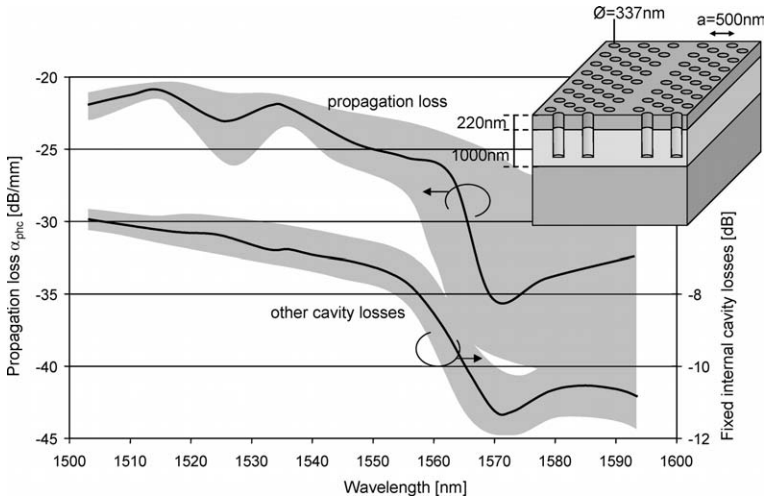

Figure 6.22: Propagation losses of a deeply etched W1 photonic crystal waveguide with a pitch of $500nm$ and a hole diameter of $337nm$. The losses are extracted by making a linear fit of the internal cavity loss τ^2 plotted in figure 6.21 as a function of waveguide length. This gives us the propagation losses as a function of wavelength, as well as the other cavity losses.

tice hole size of $330nm$, but with different defect hole sizes between $180nm$ and $220nm$. The waveguide length for the spectra shown here is $100\mu m$ in each case. We used the mode mixer to couple to both even and odd modes. For each spectrum we can observe the band edge of the asymmetric first-order guided mode, which is below the light line. The band edge is visible as a clear drop in transmission where the waveguide supports no mode. Towards the right, the rising transmission reflects the presence of the even ground mode above the light line.

As we expect, the band edge of the guided mode shifts towards short wavelengths for larger defects. Intuitively, this can be explained by the decrease in optical volume of the waveguide as high-index Silicon is replaced by low-index air.

6.5.3 Generation 2: Conclusion

Using the thick oxide has proven to reduce the substrate leakage which dominated the losses in our first-generation structures. However, the deep etching, which caused major sidewall roughness, introduced a new loss mechanism. This keeps the losses for single-mode photonic wires as high as $13dB/mm$. For photonic crystal waveguides, we have

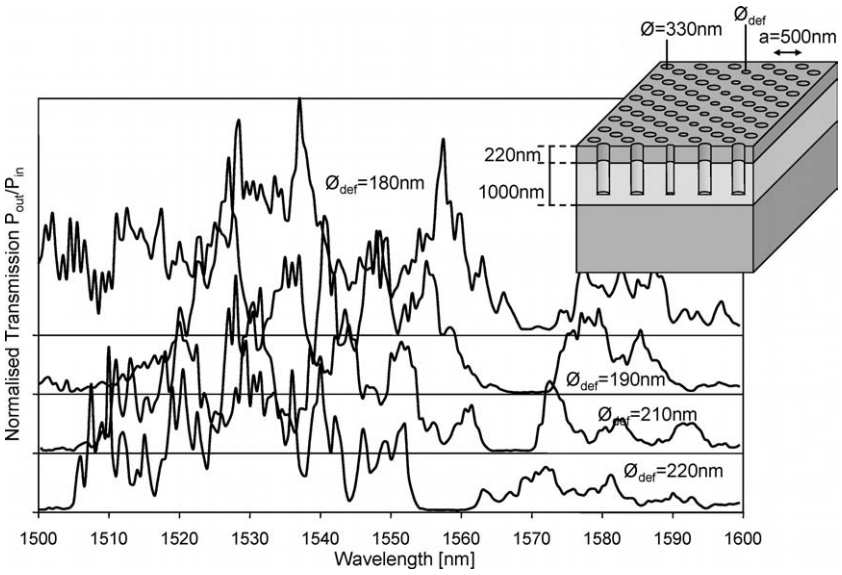

Figure 6.23: Transmission spectra of deeply etched W1 photonic crystal waveguide with small defect holes of $100\mu m$ length. The lattice pitch is $500nm$ and the hole diameter is $330nm$. We have measured 4 different defect hole sizes. For incoupling we used the mode mixer to couple light into the odd mode. We can see a dip in the transmission around $1560nm$, which coincides with the right band edge of the asymmetric first-order. On the right of the gap we see the even mode above the light line, with higher propagation losses. As the defect hole size $\emptyset_{def}$ gets larger, the bands are shifted towards lower wavelengths.

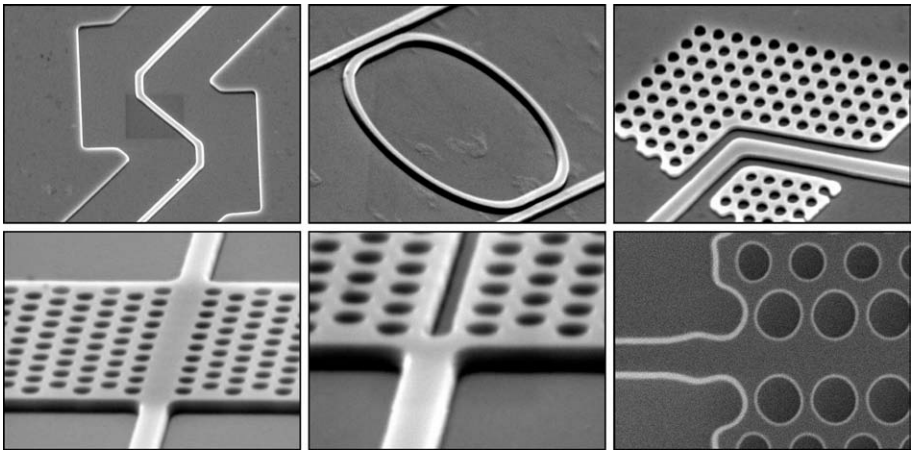

Figure 6.24: Nanophotonic structures in SOI with a thick oxide buffer and a Silicon-only etch. Top row: A wire with two corner mirrors, a racetrack resonator, a wire embedded in a photonic crystal bend. Bottom row: A W1 photonic crystal waveguide, a W1 waveguide with a single trench, a W1 waveguide with bigger border holes.

measured propagation losses of $21dB/mm$ and found otherwise good agreement between the measured spectra and the calculated band diagrams.

6.6 Generation 3: Silicon-only etch

For high-contrast photonic waveguides, deep etching into SOI seemed to be the most attractive option, because of both higher lateral and vertical contrast, and because scattering at the bottom of the photonic crystal holes was less severe. However, we have shown that sidewall roughness is a serious performance killer. Therefore, we have switched to a Silicon-only etch process. Figure 6.24 shows a number of structures fabricated with this process. We can see that the etch roughness has become barely noticable.

6.6.1 Photonic Wires

As roughness is the major cause of propagation loss in photonic wires, we expect very low propagation losses in these structures. These low losses should translate to excellent Fabry-Perot cavities between the cleaved facets. Figure 6.25 shows a detail of the transmission spectra

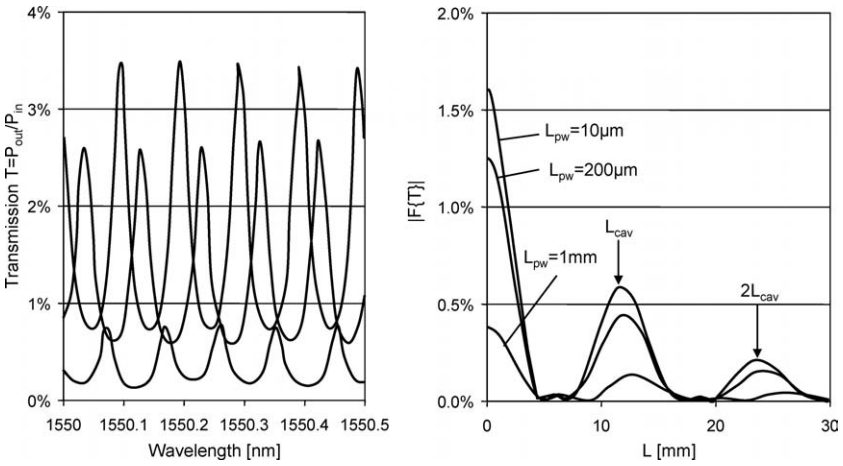

Figure 6.25: Transmission spectra of $500nm$ wide photonic wires with a Silicon-only etch. For clarity, only the wire lengths $L_{pw} = 10, 200, 1000\mu m$ are plotted. From the very pronounced Fabry-Perot spectrum we can extract the losses through the fringe contrast. Alternatively, we can Fourier transform the spectra to extract the losses from the peak ratio in the Fourier transform (right).

of $500nm$ wires with a length of $10\mu m$, $200\mu m$ and $1mm$. The Fabry-Perot fringes are clearly visible and well defined, even for the longest waveguide.

In the resulting Fourier transform, also shown in figure 6.25, we can easily discern the first and second peak corresponding to $L = L_{cav}$ and $L = 2L_{cav}$. With these clean measurements, we can easily extract the propagation losses, either from the fringe contrast or from the ratio of the peaks in the Fourier transform. Note that the peaks are broad in the Fourier transform as we only use a measurement window of $0.5nm$. Using a broader window would result in narrower peaks.

In figure 6.26, the propagation losses for photonic wires of $400nm$, $450nm$ and $500nm$ width are shown, measured using the fringe contrast method. The numerical values, as well as the error flags, are given in table 6.3. When using the peak ratio in the Fourier transform, we obtain similar results. Both results are located well within each other's error flags. For $500nm$ wires, we find propagation losses as low as $0.24dB/mm$. As the wires get narrower, the propagation loss increases rapidly.

As we have seen in chapter 4, the substrate leakage for a photonic wire with $500nm$ width is already larger than $0.1dB/mm$. This means

Width	Propagation loss	Maximum taper loss	Theo. substr. leak.
$400nm$	3.4 $\pm$ 0.17 dB/mm	-0.28 $\pm$ 0.04 dB	0.36 dB/mm
$440nm$	0.95 $\pm$ 0.18 dB/mm	-0.23 $\pm$ 0.03 dB	0.18 dB/mm
$450nm$	0.74 $\pm$ 0.09 dB/mm	-0.22 $\pm$ 0.01 dB	0.17 dB/mm
$500nm$	0.24 $\pm$ 0.16 dB/mm	-0.20 $\pm$ 0.03 dB	0.11 dB/mm

Table 6.3: Propagation losses of photonic wires as a function of wire width. The wires are etched only through the Silicon. The calculated values for the substrate leakage are also given.

that a large fraction of the propagation losses for our photonic wires are due to substrate leakage. The theoretical calculated losses due to substrate leakage are also given in table 6.3. We see that the actual propagation losse increase more rapidly than the losses due to substrate leakage.

We can now try to calculate the residual losses from the intercept of the fitted line with the ordinate axis. For the three wire widths, we assumed that the reflection R at the facets is the same, and close to the theoretical value. We could then calculate an upper limit of the losses of the tapers. These are also given in table 6.3.

6.6.2 Rings and racetracks

With these high-quality photonic wires we also fabricated ring and racetrack resonators. Figure 6.27 shows the transmission spectrum of a racetrack resonator with a radius $r_{ring} = 5\mu m$ and a coupling section of $L_{coup} = 3\mu m$. We can see clear peaks in the drop port and similar dips at the pass port. The large amount of 'noise' in the pass port is due to the Fabry-Perot formed by the cleaved facets. The cavity has a Q larger than 3000 and a free spectral range of $14nm$. At resonance, the output power in the drop port is approximately 80% of the output power in the pass port off-resonance.

We have also measured ring resonators. Because the coupling section in a ring resonator is much shorter than in a photonic wire, the coupling efficiency of the wire to the ring is much lower. This translates into both a higher Q (because the ring loses less light to the wire) and a lower overall coupling efficiency to the drop port. We have characterised ring resonators with a Q of 8000 but a coupling efficiency of only 1%.

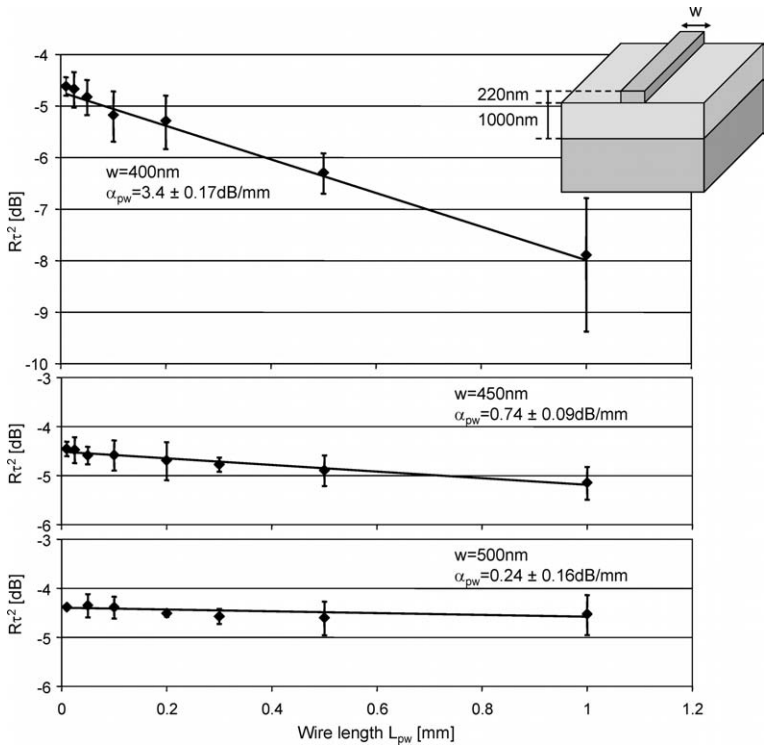

Figure 6.26: Propagation losses of photonic wires with a Silicon-only etch. The losses for each waveguide are extracted from the fringe contrast in the Fabry-Perot spectrum. The losses as function of length are plotted for photonic wires with a width of $400nm$, $450nm$ and $500nm$.

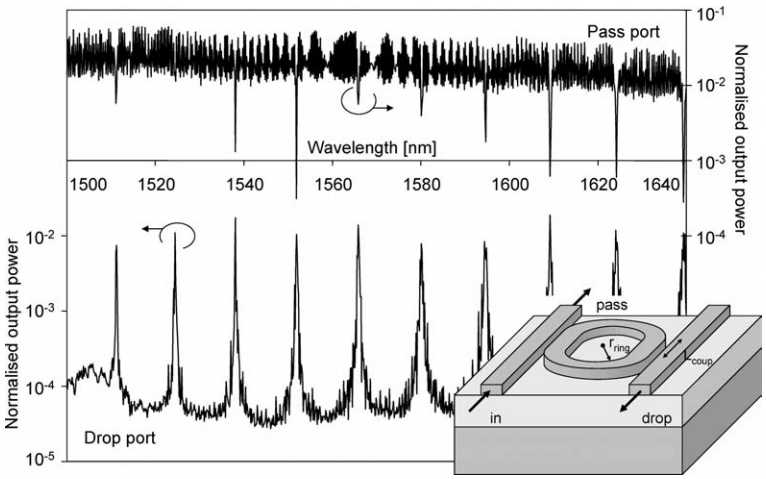

Figure 6.27: Transmission spectrum of a racetrack resonator in the pass port and the drop port. The resonator has a Q of over 3000 and a coupling efficiency at resonance of 80%.

6.6.3 Photonic Crystal Waveguides

We also measured the propagation losses of photonic crystal waveguides. Figure 6.28 shows the transmission spectra of a W1 photonic crystal waveguide with light coupled in through the even/odd mode mixer. We can see that even for very long waveguides ($L_{phc} = 1mm$) we can still measure transmission over a large part of the wavelength range.

If we now extract the measurement losses by Fourier transforming the spectra and measuring the peak height directly, we find the propagation losses plotted in figure 6.29. Around $1525nm$, the odd mode is guided and has a propagation loss as low as $7.5dB/mm$, which is much lower than the propagation losses of the deeply etched structures. Also, the other cavity losses have dropped with a few dB.

Figure 6.30 shows another photonic crystal waveguide, this time a W1 with little defect holes with a diameter $\emptyset_{def} = 200nm$. Again, we used the mode mixer for incoupling. For this structure, we have simulated the exact band structure, and we can match the regions of low propagation losses exactly to the guided modes in the band diagram. However, the lowest losses in this structure are still of the order of $40dB/mm$. This is most likely due to scattering at sidewall roughness.

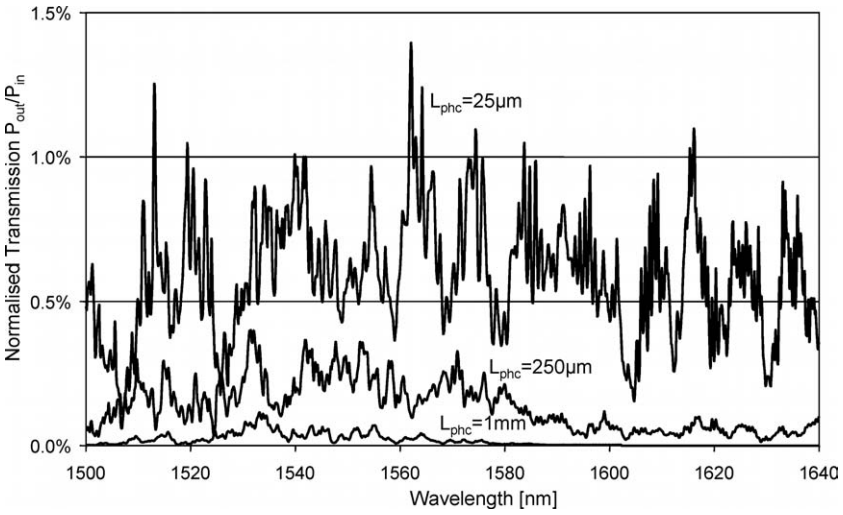

Figure 6.28: Transmission spectra of W1 photonic crystal waveguide with a Silicon-only etch. The lattice has a pitch of $500nm$ and the holes have a diameter of $320nm$. We used incoupling with the mode mixer. Note that there is still significant transmission through a waveguide of $1mm$ length

6.6.4 Generation 3: Conclusion

With the use of a Silicon-only etch we have made nanophotonic waveguides of very high quality. For the photonic wires, this translates directly into low propagation losses, as low as $0.24dB/mm$ for single mode wires of $500nm$ width. We have also demonstrated ring resonators and racetrack resonators. Photonic crystal waveguides are still considerably more lossy than photonic wires, but we have also shown $7.5dB/mm$ loss in a W1 waveguide, be it in the less useful asymmetric mode.

6.7 Generation 4: Coupling Structures

Apart from containing a number of new designs, the devices on the PICCO_3 mask are processed with two-step processing, including fibre couplers for reliable incoupling and outcoupling [3]. Also, we have fabricated compact spot-size converters based on the interference tapers

[3]The research on the grating-based fibre coupler has been carried out by Dirk Taillaert in the context of his PhD.

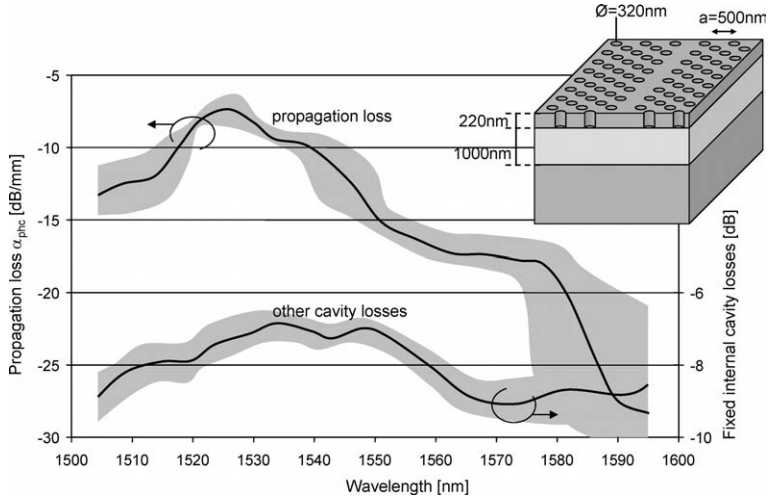

Figure 6.29: Propagation losses of a W1 photonic crystal waveguide with Silicon-only etch. The lattice has a pitch of $500nm$ and the holes a diameter of $320nm$. The losses are extracted by making a linear fit of the internal cavity loss τ^2 as a function of waveguide length. τ^2 is extracted from the peak height in a localised Fourier transform of the spectra in figure 6.28. This gives us the propagation losses as a function of wavelength, as well as the other cavity losses. Around $1525nm$, the asymmetric first-order mode has a propagation loss of approximately $7.5dB/mm$.

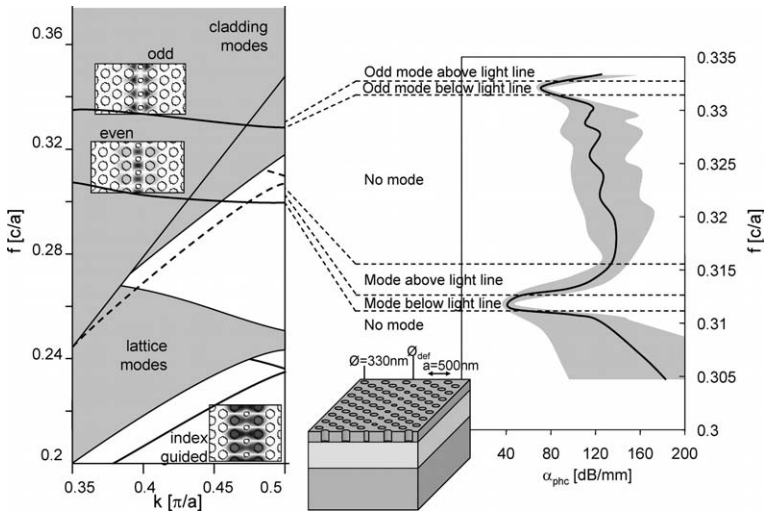

Figure 6.30: Propagation losses of a W1 waveguide with small defect holes matched to the calculated band diagram. We can see that the regions of low propagation loss match exactly the guided modes below the light line. The lattice constant is $500nm$, the hole diameter in the bulk of the lattice is $320nm$ and the defect hole diameter $\emptyset_{def} = 200nm$.

discussed in chapter 3[4]. Figure 6.31 shows some fabricated fibre couplers and spot-size converters.

6.7.1 Vertical Fibre couplers

To measure the coupling efficiency of the fibre couplers we have measured the transmission of a ridge waveguide from fibre to fibre using the fibre coupler both for incoupling and outcoupling. The $7mm$ long ridge waveguide is $10\mu m$ wide at the position of the fibre coupler and tapered down in the middle to a width of $500nm$ to filter out the higher-order waveguide modes. We normalised the measurements to the power measured when connecting the detector directly to the light source. When neglecting the losses in the waveguides, the transmission gives us the efficiency of two fibre couplers. When expressed in dB, the efficiency of a single coupler is half that of both couplers in series.

Figure 6.32 shows the coupling efficiency of a single fibre coupler with a pitch of $610nm$ and a trench width of $305nm$. The curve maxi-

[4]The research on the integrated interference-based spot-size converter has been carried out by Bert Luyssaert in the context of his PhD.

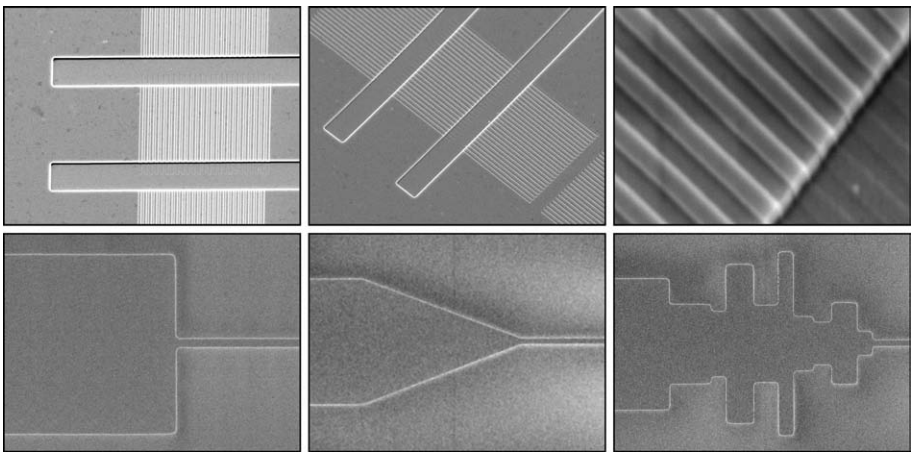

Figure 6.31: Fibre couplers and compact spot-size converters. Top row: 1-D fibre couplers, with a shallow etch and a waveguide trench with a deep etch. Bottom row: Butt-coupled waveguides, a short linear taper and an interference-based spot-size converter.

mum is at $-6.8dB$, which corresponds to a coupling efficiency of 21%. We can also see that the $3dB$ bandwidth is $60nm$.

While the coupling from the fibre is predominantly to the ground mode of the broad ridge waveguide, there is some residual coupling to the higher-order waveguide modes. These higher-order modes should be filtered out using a mode filter (in our case, an adiabatic taper down to $500nm$ width and back). If not, the different waveguides mode can either add up or not at the outcoupling facet and cause a wavelength dependent modulation of the output power (so-called mode beating).

6.7.2 Compact Spot-size Converters

Using the fibre couplers, we have also characterised various compact spot-size converters. Figure 6.33 compares the output from four different coupling structures between a $10\mu m$ wide ridge and a $560nm$ wire: The reference fibre couplers with an adiabatic mode filter, a short linear taper (only $50\mu m$ long), two butt-coupled waveguides and an optimised 10-section interference coupler with a length of $15.4\mu m$.

We can see that the $15.4\mu m$ long interference coupler performs with similar efficiency as the three times longer linear taper. When subtracting the transmission from the reference curve, we find that a single cou-

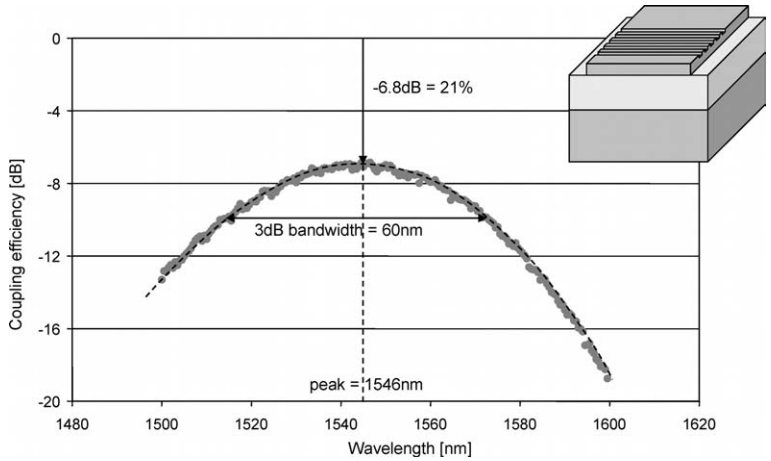

Figure 6.32: Coupling efficiency of a 1-D fibre coupler. The curve is a Gaussian with a maximum coupling efficiency of $-6.8dB$ or 21%, and a $3dB$ bandwidth of $60nm$.

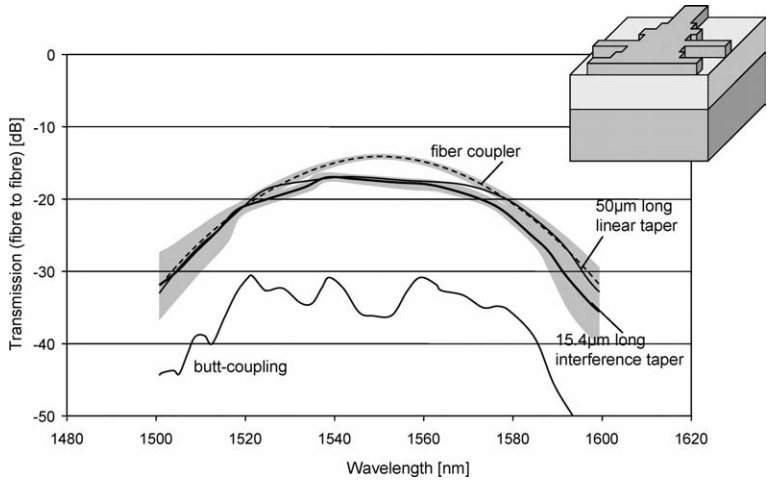

Figure 6.33: Transmission of compact spot-size converters. The graph compares the fibre-to-fibre transmission of a $7mm$ long structure consisting a $10\mu m$ ridge waveguide with a spot-size converter down to $560nm$ and back. Incoupling and outcoupling is done using the vertical fibre couplers. The top reference curve is the fibre coupler with an adiabatic tapering. The other structures are illustrated in figure 6.31: a $50\mu m$ long linear taper, a butt-coupling between the broad and narrow waveguide, and an optimised 10-section interference coupler.

	$500nm$ wire	PhC WG (SOI)	PhC WG (membrane)
Deep etch	34 dB/mm	21 dB/mm	N/A
Si-only etch	0.24 dB/mm	7.5 dB/mm	N/A
Literature	0.35 dB/mm [1]	2.4 dB/mm [1] 2 dB/mm [2]	0.7 dB/mm [1] 0.7 dB/mm [2] 0.7 dB/mm [3]

[1] Vlasov, IBM: W1, E-beam lithography in SOI or Si membrane [48]
[2] Notomi, NTT: W0.7, E-beam lithography in SOI or Si Membrane[14]
[3] Sugimoto, FESTA: W1, E-beam lithography in GaAs membrane

Table 6.4: Propagation losses of nanophotonic waveguides compared.

pler has an efficiency of $-1.5dB$ (70%) or better over a bandwidth of more than $50nm$.

6.8 Conclusion

In the time span of this work, we have characterised a large number of components, fabricated with a number of different fabrication processes. In the end, we found that the best process consisted of a Silicon-only etch, resulting in low sidewall roughness and therefore low propagation losses. The results for these structures are listed again in table 6.4, and compared to the state-of-the-art e-beam fabricated structures published in literature [48, 14], both in SOI and membranes in air. For our structures, fabricated with deep UV lithography, we have measured propagation losses as low a $0.24dB/mm$ for photonic wires and $7.5dB/mm$ for photonic crystals.

We have also demonstrated vertical fibre-coupler structures with a coupling efficiency of over 20%, as well as very short, broad-band spotsize converters with a efficiency of 70%.

One of the limitations of our measurements was the light source. It has a limited bandwidth of $140nm$, which is often not enough to span the photonic bandgap of a photonic crystal lattice. Also,the emission linewidth is very narrow, which makes it impossible to average the transmission of a Fabry-Perot spectrum without scanning the whole range with sufficient accuracy to resolve the individual peaks. Using a broad-band light source in combination with a spectrum analyser would facilitate the measurements of nanophotonic waveguides considerably.

Chapter 7

Conclusions

"Well Ug, I see photonic crystals definitely made an impact."

Chapter 7

Conclusions

On the road towards large-scale integration of photonic structures nanophotonics marks the way. In this work we have discussed both the theory and the application of nanophotonic waveguide structures. For this, we designed, fabricated and characterised nanophotonic components in Silicon-on-insulator.

7.1 Nanophotonics and Photonic Crystals

The key aspect of nanophotonics is the use of wavelength-scale structures with a very high refractive index contrast. The two most prominent building blocks for integrated nanophotonic circuits are *photonic wires* and *photonic crystals*. Photonic wires function like conventional waveguides, be it on a smaller scale. Photonic crystals, as we have seen in chapter 2, are periodic structures with a high index contrast. When properly designed, they have a photonic bandgap: a wavelength range in which light cannot propagate through the structure. This unique aspect of photonic crystals makes it possible to confine light inside a small region, creating cavities and waveguides.

While photonic crystals demonstrate the most diverse and interesting properties, we have shown in chapter 3 that it is quite a challenge to design a good nanophotonic waveguide with them. Photonic wires, with their simple index guiding mechanism, are a much better candidate for pure waveguiding and interconnection. With their high index contrast, it is possible to make sharp bends, corner mirrors, and even wavelength-selective elements in the form of ring resonators.

On the other hand, the peculiar properties of photonic crystals can be used for a number of compact, wavelength-selective functions. This can be done by tuning intentional defects, or by changing the properties of the periodic lattice itself.

One of the key requirements for the integration of nanophotonic waveguides into larger circuits is the control of the waveguide propagation losses. In chapter 4 we have discussed the various loss mechanisms in both photonic wires and photonic crystal waveguides, and a number of design rules to keep them in check. An important aspect is the choice of the material system. We have shown that layer structures with a high vertical index contrast are beneficial to reduce both scattering inside the photonic crystal holes (intrinsic losses) and scattering at sidewall roughness. Silicon-on-insulator, the material we used for the fabrication of our structures, falls into this category. Layer structures with a low vertical index contrast, like $III - V$ semiconductor heterostructures, are much more prone to scattering at sidewall roughness.

7.2 Fabrication with Deep UV Lithography

For the fabrication we had access to the advanced processing facilities of IMEC in Leuven. Instead of using e-beam lithography, which is used for the majority of nanophotonic research worldwide, we used $248nm$ deep UV lithography. This mass-manufacturing technique used for high-end CMOS devices offers a road towards commercialisation of nanophotonics. Silicon-on-insulator proved to be an ideal material, as it is compatible with the Silicon-based processing as well as a very good material for nanophotonic circuitry.

However, as described in chapter 5, the transition from CMOS processes to an optimised fabrication process for nanophotonic structures was not straightforward. Because of the tight requirements for alignment and feature dimensions in a nanophotonic circuit, very careful process characterisation is needed in order to determine the necessary biases and optical proximity corrections during lithography.

For the etching of the structures into the SOI, we explored a number of routes. By far the best results were achieved by etching only the top Silicon layer, in combination with a plasma treatment of the developed photoresist. This process introduces very little sidewall roughness.

The optimised process is capable of fabricating a large variety of nanophotonic structures in SOI, using $248nm$ deep UV lithography as

the pattern definition technique. However, a number of improvements are possible. We have shown that the use of better resists can increase the process window, making the process more reproducible. The same is true for advanced mask technology, like phase shifting masks.

Even though the sidewalls are already quite even, further smoothening is possible by oxidation, and oxide deposition can make the SOI layer structure vertically symmetric.

7.3 Characterisation

For the characterisation of the fabricated structures, we coupled light into the structures and measured the transmission at the other end. In chapter 6, the various measurement and analysis techniques are discussed in detail. As expected, photonic wires proved to be the best nanophotonic waveguides, once sidewall roughness was under control. Propagation losses are as low a $0.24dB/mm$ for $500nm$ wide wires, and we measured ring resonators with a quality factor Q of 8000.

The photonic crystal waveguides have considerably higher propagation losses. The lowest measured value of $7.5dB/mm$ is still an order of magnitude larger than the losses in photonic wires. However, while the losses are still somewhat high, we found excellent agreement between the calculated dispersion diagrams (or band diagrams) and the transmission spectra.

Apart from nanophotonic waveguides, we have also fabricated coupling structures to interface a nanophotonic circuit with single-mode fibres used in telecommunications. The grating-based fibre couplers have an experimental efficiency of more than 20%. These fibre couplers couple the light from the fibre into a $10\mu m$ wide ridge waveguide in Silicon-on-insulator. To further reduce the spot size to that of a nanophotonic waveguide, we demonstrated an interference-based spot-size conversion with an efficiency of over 70%.

One of the main difficulties encountered in characterising nanophotonic structures is the required measurement resolution. Because the structure typically consists of a large number of cavities, a very fine wavelength scan is needed, while at the same time a broad wavelength range is desired to map out the structure's wavelength-dependent behaviour. The use of a broadband source instead of a tunable laser might facilitate this.

7.4 Conclusion

To summarise this work, we can state that we have demonstrated a viable route towards the commercial adoption of nanophotonics for large-scale integration. Photonic wires with acceptably low losses and photonic crystals with strong wavelength-dependent behaviour offer a wide range of functional elements. Fibre couplers and spot-size converters with a small footprint take care of the interface with the outside world, and make wafer-scale testing a possibility. Fabrication based on mature CMOS technology has proven it can provide the large-scale manufacturing base for nanophotonic components in SOI. In the following years, we will witness the evolution of photonics akin to that of electronics, and even the integration of both into a hybrid opto-electronic circuit. Already, the industry is looking to photonics for on-chip communications. Low-loss nanophotonic waveguides, as we have demonstrated here, might be one of the ways to make this happen.

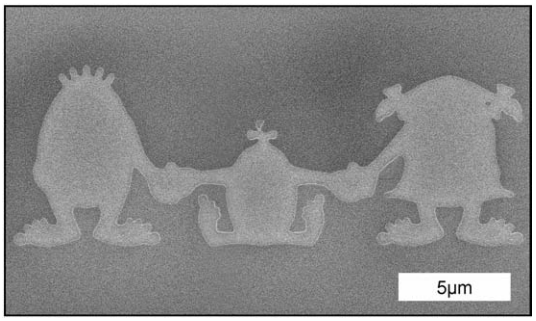

Appendix A

Deep UV Masks

In this appendix we will give an overview of the contents of the four PICCO masks. The useful mask area is approximately $2.8 \times 2.6 cm^2$ at $1X$. Because at this size we can only print seven columns of dies across a $200mm$ wafer, we have subdivided the mask into cells that can be processed independently. This is also useful as we can group the structures that require different processing conditions. During illumination, the stepper can select a specific cell from the mask by choosing a rectangular aperture. The wafer layout is also dependent on the dimensions of the individual cells.

A.1 PICCO_1

The PICCO_1 mask was conceived as a follow-up for the fabrication tests performed with the in-house masks at IMEC. However, some existing photonic crystal waveguide designs from the literature were included. Because of this limited availability of designs, the mask is only partially filled. The PICCO_1 mask contains 4 cells:

- **CLEAVE**: Large, cleavable lattices to evaluate etch tests in cross section. (Design: Wim Bogaerts, Dirk Taillaert)
- **OPC**: Structures with large varieties of hole configurations in photonc crystal waveguides, bends and splitters to characterise optical proximity effects and determine the necessary corrections. (Design: Wim Bogaerts, Dirk Taillaert)
- **FIBCOUP**: First designs of the out-of-plane fibre coupler. (Design: Dirk Taillaert)

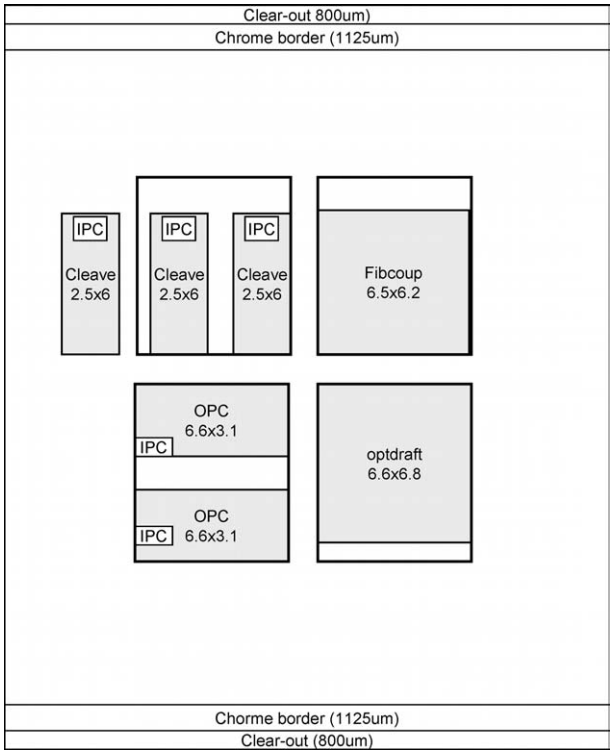

Figure A.1: The PICCO_1 mask layout. All dimensions are in millimeters at $1X$.

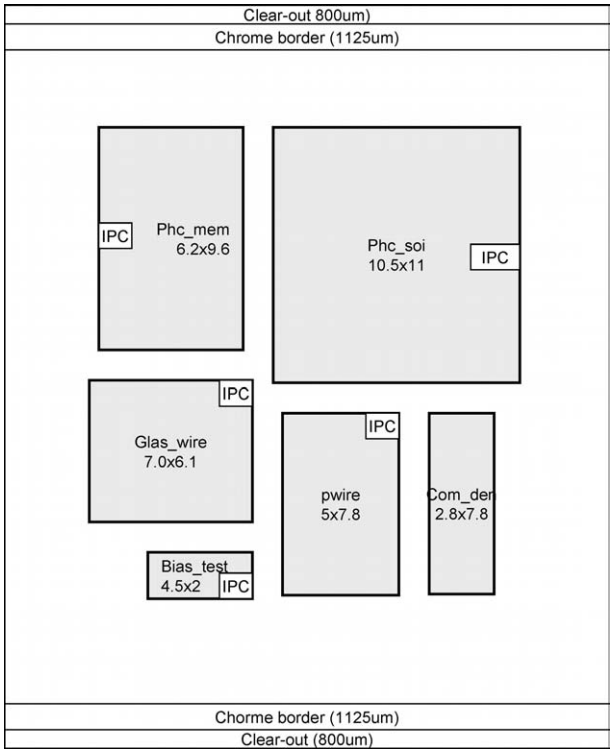

Figure A.2: The PICCO_2 mask layout. All dimensions are in millimeters at $1X$.

- **OPTDRAFT**: Photonic crystal and photonic wire designs from literature, adapted to the right SOI layer structure. (Design: Wim Bogaerts, Dirk Taillaert)

A.2 PICCO_2

Even though structures based on the PICCO_1 mask had not been fabricated, improved designs motivated the fabrication of another mask. PICCO_2 contains 6 cells:

- **PWIRE**: Photonic wire waveguides, ring resonators and sharp bends. (Design: Pieter Dumon, Wim Bogaerts)

- **PHC_SOI**: Photonic crystal waveguides for the SOI layer stack. (Design: Joris Van Campenhout, Wim Bogaerts)

- **PHC_MEM**: Photonic crystal waveguide, optimised for Silicon membranes. (Design: Joris Van Campenhout, Wim Bogaerts)
- **COM_DEN**: Photonic crystal waveguide designs contributed by Research Center COM, Denmark. (Design: Morten Thorhauge)
- **GLAS_WIRE**: Photonic wire designs from Glasgow University, UK. (Design: Harold Chong)
- **BIAS_TEST**: Advanced metrology structures to establish necessary bias to fabricate different types of structures together (dense holes, isolated lines,) (Design: Wim Bogaerts)

A.3 PICCO_3

The PICCO_3 mask introduces 2 novel aspects: We standardised the procedure for contributions from outside groups through a *Deep UV processing collaboration*. Also, we included devices for 2-step processing: First, the fibre couplers would be processed, and then the deeply etched structures (see chapter 5. For this, we stipulated standard positions for the incoupling waveguides. The PICCO_3 mask contains 7 cells:

- **A1**: Photonic crystal waveguides. This cell can be combined with the fibre couplers of cell B. (Design: Pieter Dumon, Wim Bogaerts)
- **A2**: Compact spot-size converters and 2-D fibre couplers. This cell can be combined with the fibre couplers of cell B. (Design: Bert Luyssaert)
- **A3**: Experimental designs for vertical directional couplers. This cell can be combined with the fibre couplers of cell B. (Design: Gunther Roelkens, Wim Bogaerts)
- **B**: Fibre couplers for input and output matching the structuren on the cells A1 through A3. (Design: Dirk Taillaert, Wim Bogaerts)
- **LCPIC**: Photonic wire based Mach-Zehnder circuits for experiments with a liquid crystal top cladding on SOI waveguides. (Design: Hans Desmet, Wim Bogaerts)
- **DUVPC**: Combined contributions from the contributors through the deep UV processing collaboration: Max-Planck Institut Hall

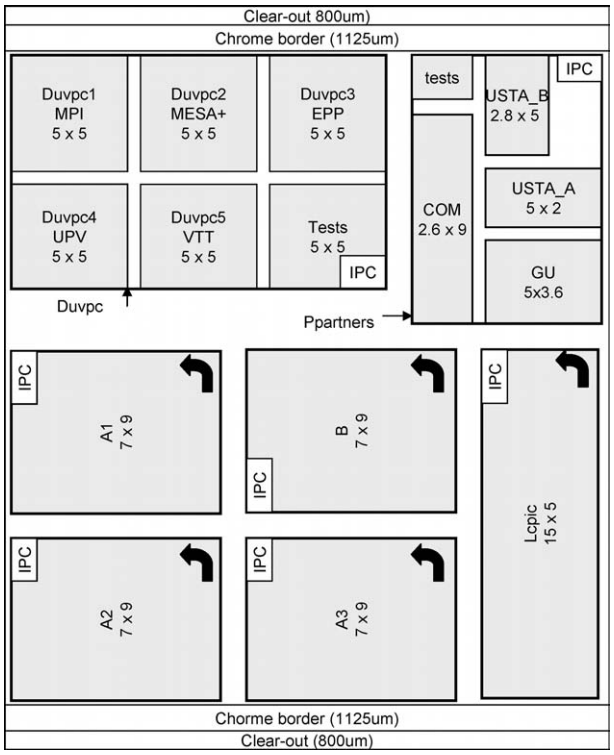

Figure A.3: The PICCO_3 mask layout. All dimensions are in millimeters at $1X$. Some designs are rotated 90 degrees and are marked with a bent arrow.

(Germany), MESA+, Twente (The Netherlands), Ecole Polytechnique de Palaiseau (France), the Technical University of Valencia (Spain) and VTT Technical Research Centre (Finland).

- **PPARTNERS**: Combined contributions from the partners from the IST-PICCO project: Glasgow University (UK), Research Center COM (Denmark) and University of Saint-Andrews (UK).

A.4 PICCO_4

PICCO_4 builds further upon PICCO_2 and PICCO_3 designs, and also includes designs from PICCO and external partners. Also, designs with the 2-D fibre couplers for 2-step processing have been added, as well as a rudimentary polarisation-diversity circuit. At the design time, no improved photonic crystal waveguide designs were available in IMEC. PICCO_4 contains 9 cells:

- **A4**: Advanced photonic wire based directional couplers for heterogeneous integration through wafer bonding. This cell can be combined with the fibre couplers of cell B on PICCO_3. (Design: Gunther Roelkens, Pieter Dumon)

- **A5**: Advanced designs for compact spot-size converters, and new designs for photonic crystal waveguides. This cell can be combined with the fibre couplers of cell B on PICCO_3. (Design: Bert Luyssaert, Pieter Dumon)

- **A6**: Advanced photonic wire based components for WDM applications. This cell can be combined with the fibre couplers of cell B on PICCO_3. (Design: Pieter Dumon, Wim Bogaerts)

- **A7**: Very compact, photonic wire based AWGs and periodically varied photonic wires for nonlinear functions. This cell can be combined with the fibre couplers of cell B on PICCO_3. (Design: Pieter Dumon, Dries Van Thourhout, Gino Priem, Wim Bogaerts)

- **SENSOR_A**: A grating-based two-dimensional stress-sensor and a simple polarisation-diversity circuit with ring resonators mathcing the 2-D fibre couplers on the cell SENSOR_B. (Design: Pieter Dumon, Reinhard Laroy, Dirk Taillaert, Wim Bogaerts)

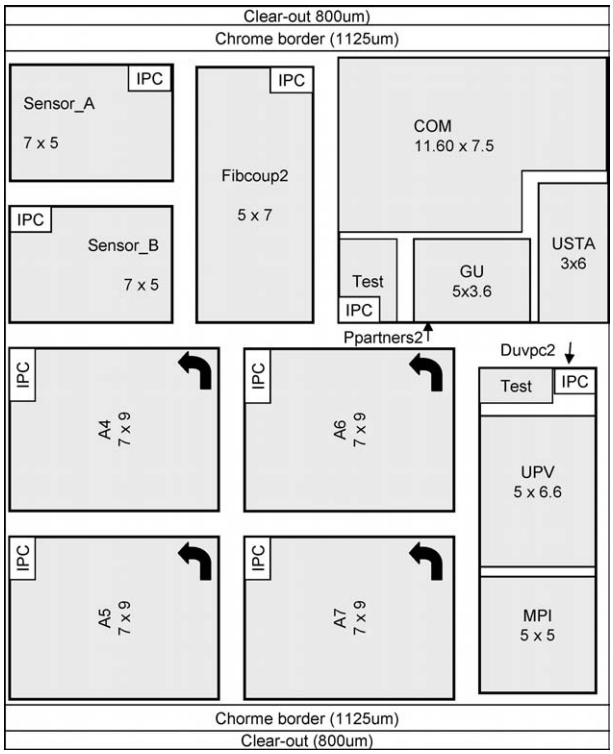

Figure A.4: The PICCO_4 mask layout. All dimensions are in millimeters at $1X$. Some designs are rotated 90 degrees and are marked with a bent arrow.

- **SENSOR_B**: Gratings and 2-D fibre couplers for 2-step processing in combination with SENSOR_A. (Design: Pieter Dumon, Reinhard Laroy, Dirk Taillaert, Wim Bogaerts)
- **FIBCOUP2**: Advanced designs for 1-D and 2-D fibre couplers. (Design: Dirk Tailaert)
- **DUVPC2**: Combined contributions from the external contributors through the deep UV processing collaboration: Max-Planck Institut Hall (Germany) and the Technical University of Valencia (Spain).
- **PPARTNERS_2**: Combined contributions from PICCO partners: Glasgow University (UK), Research Center COM (Denmark) and University of Saint-Andrews (UK).

Appendix B

Publications

B.1 International Journals

1. 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, *Photon. Technol. Lett.*, accepted for publication.

2. W. Bogaerts, R. Baets, P. Dumon, V. Wiaux, S. Beckx, J. Van Campenhout, D. Taillaert, B. Luyssaert, P. Bienstman, D. Van Thourhout, Nanophotonic Waveguide Structures in Silicon-on-Insulator (invited), submitted to *J. Lightwave Technol.*.

3. W. Bogaerts, R. Baets, P. Dumon, V. Wiaux, S. Beckx, J. Van Campenhout, D. Taillaert, B. Luyssaert, P. Bienstman, D. Van Thourhout, Basic Structures for Photonic Integrated Circuits in Silicon-on-Insulator (invited), *Opt. Exp.*, accepted for publication.

4. P. Sanchis, J. Garcia, J. Martí, W. Bogaerts, P. Dumon, D. Taillaert, R. Baets, V. Wiaux, J. Wouters, S. Beckx, Experimental demonstration of high efficiency coupling technique for photonic crystal circuits, submitted to *Photon. Technol. Lett.*.

5. L. H. Frandsen, P. I. Borel, Y. X. Zhuang, A. Harpoth, M. Thorhauge, M. Kristensen, W. Bogaerts, P. Dumon, R. Baets, V. Wiaux, J. Wouters, S. Beckx, Ultra Low Loss 3 dB Photonic Crystal Waveguide Splitter, submitted to *Opt. Lett.*.

6. P. Dumon, W. Bogaerts, V. Wiaux, J. Wouters, S. Beckx, J. Van Campenhout, D. Taillaert, B. Luyssaert, P. Bienstman, D. Van Thourhout and R. Baets, Low-loss SOI Photonic Wires and Ring Resonators Fabricated with Deep UV Lithography, submitted to *Photon. Technol. Lett.*

7. W. Bogaerts, P. Bienstman and R. Baets, Scattering at sidewall roughness in photonic crystal slabs, *Opt. Lett.*, 28(9): 689-691, Apr. 2003.

8. 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. Top. Quant. Electron.*, 8(4): 928-934, Oct. 2002.

9. D. Taillaert, W. Bogaerts, P. Bienstman, T.F. Krauss, P. Van Daele, I. Moerman, S. Verstuyft, K. De Mesel and R. Baets, An Out-of-Plane Grating Coupler for Efficient Butt-Coupling Between Compact Planar Waveguides and Single-Mode Fibers, *IEEE J. Quant. Electron.*, 38(7): 949-955, Jul. 2002.

10. W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets and D. De Zutter, Out-of-plane scattering in 1-D photonic crystal slabs, *Opt. Quant. Electron.*, 34(1/3), Kluwer Academic Publishers, : 195-203, Mar. 2002.

11. W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets and D. De Zutter, Out-of-Plane Scattering in Photonic Crystal Slabs, *IEEE Phot. Technol. Lett.*, 13(6): 565-567, Jun. 2001.

B.2 International Conference Proceedings

12. R. Baets, P. Bienstman, W. Bogaerts, P. Dumon, D. Taillaert, B. Luyssaert, Nanophotonics based on high-index-contrast photonics devices (invited), OECC/COIN, Yokohama, Japan, Jul. 2004, accepted for publication.

13. 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*, Strasbourg, France, Apr.2004, accepted for publication.

14. R. Baets and W. Bogaerts, Nanophotonic waveguides in SOI by means of deep UV lithography (invited), *PECS-V*, Kyoto, Japan: 160, Mar. 2004.

15. B. Luyssaert, P. Vandersteegen, W. Bogaerts, V. Wiaux, J. Wouters, S. Beckx, R. Baets, Compact Photonic Spot-Size COnverter, *PECS-V*, Kyoto, Japan: 193, Mar. 2004.

16. M. Kristensen, L.H. Frandsen, A. Harpoeth, Y.X. Zhuang, P.I. Borel, W. Bogaerts, P. Dumon, R. Baets, V. Wiaux, J. Wouters, S. Beckx, Large bandwidth 60° bend in photonic crystal waveguides, *PECS-V*, Kyoto, Japan: 93, Mar. 2004.

17. L.H. Frandsen, P.I. Borel, A. Harpoeth, M. Kristensen, M. Thorhauge, W. Bogaerts, P. Dumon, R. Baets, V. Wiaux, J. Wouters, S. Beckx, "Ultra low loss photonic crystal waveguide 3 dB splitter", *PECS-V*, Kyoto, Japan: 146, Mar. 2004.

18. P. Sanchis, J. Garcia, F. Cuesta-Costo, A. Martinez, J. Blasco, J. Marti, W. Bogaerts, P. Dumon, D. Taillaert, R. Baets, "Experimental demonstration of a high efficiency coupling technique for planar photonic crystal circuits", *PECS-V*, Kyoto, Japan: 128, Mar. 2004.

19. R. Baets, W. Bogaerts, D. Taillaert, P. Dumon, P. Bienstman, D. Van Thourhout, J. Van Campenhout, V. Wiaux, J. Wouters and S. Beckx, Low Loss Nanophotonic Waveguides and Ring Resonators in Silicon-on-Insulator (invited), *39th course: microresonators as building blocks for VLSI photonics*, IAP, Erice, Italy, accepted for publication.

20. Y.X. Zhuang, L.H. Frandsen, A. Harpth, M. Thorhauge, M. Kristensen, P.I. Borel, W. Bogaerts and R. Baets, Wavelength tuning of photonic crystal waveguides fabricated using 248-nm deep UV lithography , *OFC 2004*, United States: ThR2, Feb. 2004.

21. P. Dumon, W. Bogaerts, J. Van Campenhout, V. Wiaux, J. Wouters, S. Beckx and R. Baets, Low-loss photonic wires and compact ring resonators in silicon-on-insulator, *IEEE/LEOS Benelux Annual Symposium*, Enschede, Netherlands, Nov. 2003.

22. D. Taillaert, W. Bogaerts and R. Baets, Efficient coupling between submicron SOI-waveguides and single-mode fibers, *IEEE/LEOS Benelux Annual Symposium*, Enschede, Netherlands, Nov. 2003.

23. W. Bogaerts, P. Dumon, J. Van Campenhout, V. Wiaux, J. Wouters, S. Beckx, D. Taillaert, B. Luyssaert, D. Van Thourhout and R. Baets, Deep UV Lithography for Planar Photonic Crystal Structures (invited), *LEOS Annual Meeting*, 2, Tuscon, United States: 754-755, Oct. 2003.

24. P. Dumon, W. Bogaerts, J. Van Campenhout, V. Wiaux, J. Wouters, S. Beckx and R. Baets, Low-loss, single-mode photonic wires ring resonators in Silicon-on-Insulator, *LEOS Annual Meeting*, Tuscon, United States, 1: 289-299, Oct. 2003.

25. 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, *European Conference on Optical Communication (ECOC)*, Rimini, Italy, 3: 430-431, Sept. 2003.

26. W. Bogaerts, P. Dumon, V. Wiaux, J. Wouters, S. Beckx and R. Baets, Tolerance control for photonic crystal structures fabricated with deep UV lithography, *European Conference on Optical Communication (ECOC)*, Rimini, Italy, 5: 46-47, Sep. 2003.

27. W. Bogaerts, V. Wiaux, P. Dumon, D. Taillaert, J. Wouters, S. Beckx, J. Van Campenhout, B. Luyssaert, D. Van Thourhout and R. Baets, Large-scale production techniques for photonic nanostructures (invited), *Proc. SPIE*, SPIE, San Diego, United States, 5225: 101-112, Aug. 2003.

28. R. Baets, P. Bienstman, W. Bogaerts, D. Taillaert and B. Luyssaert, Coupling issues in strongly confined nanophotonic waveguides (invited), *Digest of the LEOS Summer Topical Meetings*, ISBN 0-7803-7982-9, Vancouver, Canada: 41-42, Jul. 2003.

29. R. Baets, P. Bienstman, W. Bogaerts, D. Taillaert, B. Luyssaert, Coupling issues in strongly confined nanophotonic waveguides (invited), *International Conference on Transparent Optical Networks (ICTON)*, Warsaw, Poland: 191-192, Jul. 2003.

30. W. Bogaerts, P. Bienstman and R. Baets, Sidewall Roughness in Photonic Crystal Slabs: A Comparison of High-Contrast Membranes and Low-Contrast III-V Epitaxial Structures, *European Conference on Integrated Optics (ECIO)*, Prague, Czech Republic: 349-352, Apr. 2003.

31. T.F. Krauss, R. Wilson, R. Baets, W. Bogaerts, M. Kristensen, P.I. Borel, L.H. Frandsen, M. Thorhauge, B. Tromborg, A. Lavrinenko, R.M. De La Rue, H. Chong, L. Socci, M. Midrio and D. Gallagher, Photonic Integrated Circuits using Crystal Optics (PICCO) (invited), *European Conference on Integrated Optics (ECIO)*, Prague, Czech Republic: 113-117, Apr. 2003.

32. W. Bogaerts, P. Bienstman and R. Baets, Out-of-plane scattering at sidewall roughness in photonic crystal slabs, *IEEE/LEOS Benelux Annual Symposium* , Amsterdam, Netherlands: 79-82, Dec. 2002.

33. W. Bogaerts, V. Wiaux, D. Taillaert, S. Beckx and R. Baets, SOI Photonic Crystal components fabricated with deep UV lithography, *International Workshop on Photonic and Electromagnetic Crystal Structures (PECS-IV)*, Los Angeles, United States:40, Oct. 2002.

34. D. Taillaert, W. Bogaerts, P. Bienstman and R. Baets, Out-of-plane fiber coupler for coupling to high-index-contrast waveguides, *International Workshop on Photonic and Electromagnetic Crystal Structures (PECS-IV)*, Los Angeles, United States: 155, Oct. 2002.

35. R. Baets, W. Bogaerts, D. Delbeke, P. Bienstman, D. Taillaert, and B. Luyssaert, Photonic crystal and light extraction enhancement possibilities (invited), *11th International Workshop on Inorganic and Organic Electroluminescence & 2002 International Conference on the Science and Technology of Emissive Displays and Lighting (EL 2002)*, Ghent, Belgium: 559-563, Sep. 2002.

36. W. Bogaerts, V. Wiaux, D. Taillaert, S. Beckx and R. Baets, Photonic Crystal Waveguides in SOI fabricated with deep UV lithography, *European Conference on Optical Communication (ECOC)*, Copenhagen, Denmark: 4.4.5, Sep. 2002.

37. W. Bogaerts, D. Taillaert, R. Baets, V. Wiaux and S. Beckx, Deep UV lithography for mass-manufacturing photonic crystal based components, *Workshop 2D photonic crystals*, Ascona, Switzerland: III-24, Aug. 2002.

38. W. Bogaerts, D. Taillaert, R. Baets, V. Wiaux and S. Beckx, Large-scale production techniques for photonic nanostructures using excimer laser replication (invited), *Integrated Photonics Research (IPR)*, Vancouver, Canada: IFH4, Jul. 2002.

39. W. Bogaerts, V. Wiaux, D. Taillaert, S. Beckx and R. Baets, Fabrication of Ultra-Compact Photonic Structures in Silicon-on-Insulator (SOI) Using 248 nm Deep UV Lithography, *International Conference on Transparent Optical Networks (ICTON)*, ISBN 0-7803-7375-8, Warsaw, Poland, 2: 30-33, Apr. 2002.

40. D. Taillaert, W. Bogaerts, P. Bienstman, T.F. Krauss, P. Van Daele, I. Moerman and R. Baets, An out-of-plane grating coupler for efficient butt-coupling between compact planar waveguides and single-mode fibres, *IEEE/LEOS BeneluxAnnual Symposium*, Brussels, Belgium : 17-20, Dec. 2001.

41. D. Taillaert, W. Bogaerts, P. Bienstman, T.F. Krauss, I. Moerman, P. Van Daele and R. Baets, A high-efficiency out-of-plane fibre coupler for coupling to high index contrast waveguides, *European Conference on Optical Communication (ECOC)*, Amsterdam, Netherlands: 30-31, Oct. 2001.

42. R. Baets, P. Bienstman, W. Bogaerts and D. Taillaert, Guided wave photonic ICs: photonic crystals versus total internal reflection (invited), *International Workshop on Photonic and Electromagnetic Crystal Structures (PECS-III)*, St. Andrews, United Kingdom, Jun. 2001.

43. W. Bogaerts, D. Taillaert, R. Baets, V. Wiaux, S. Beckx, M. Maenhoudt and G. Vandenberghe, SOI Photonic Crystal Fabrication Using Deep UV Lithography, *International Workshop on Photonic and Electromagnetic Crystal Structures (PECS-III)*, St. Andrews, United Kingdom, Jun. 2001.

44. D. Taillaert, W. Bogaerts, P. Bienstman, D. De Zutter and R. Baets, Design of an out-of-plane coupler for efficient butt-coupling from photonic crystal waveguides to single-mode fibres, *International Workshop on Photonic and Electromagnetic Crystal Structures (PECS-III)*, St. Andrews, United Kingdom, Jun. 2001.

45. W. Bogaerts, P. Bienstman, D. Taillaert and R. Baets, Out-of-plane scattering in photonic crystal slabs, *LEOS Annual Meeting*, Puerto Rico, Costa Rica, PD1.5, Nov. 2000.

46. W. Bogaerts, P. Bienstman, D. Taillaert and R. Baets, Out-of-plane scattering losses in 1D photonic crystal slabs, *IEEE/LEOS Benelux Annual Symposium*, Delft, Netherlands: 35-37, Oct. 2000.

47. W. Bogaerts, P. Bienstman, D. Taillaert and R. Baets, Out-of-plane scattering in 1D photonic crystal slabs, *Summer School and European Optical Society Topical Meeting in Semiconductor Microcavity Photonics*, Centro Stefano Franscini, Switzerland, Oct. 2000.

48. W. Bogaerts, D. Taillaert, R. Baets and D. De Zutter, Simulation of out-of-plane scattering losses in 2D photonic crystal waveguides, *International Workshop on Photonic and Electromagnetic Crystal Structures (PECS)*, Sendai,Miyagi, Japan, Mar. 2000.

Figures

Tables

Bibliography

[1] E.W Van Stryland, J.M. Enoch, and M. Wolfe W.L., Bass (editors). *Handbook of Optics*. OSA, 2000.

[2] E. Yablonovitch. Inhibited spontaneous emission in solid state physics and electronics. *Phys. Rev. Lett.*, 58:2059, 1987.

[3] E. Yablonovitch. Photonic band-gap structures. *J. Opt. Soc. Am. B*, 10:283–296, 1993.

[4] J. D. Joannopolous, R.D. Meade, and J.N. Winn. *Photonic Crystals - Molding the Flow of Light*. Princeton University Press, Princeton, N.J., 1995.

[5] T.F. Krauss, R.M. De La Rue, and S. Brand. Two-dimensional photonic-bandgap structures operating at near infrared wavelengths. *Nature*, 383:699–702, 1996.

[6] T.F. Krauss, B. Vogele, C.R. Stanley, and R.M. De la Rue. Use of guided spontaneous emission of a semiconductor to probe the optical properties of two-dimensional photonic crystals. *IEEE Phot. Technol. Lett.*, 71:738–740, 1997.

[7] M. Loncar, T. Doll, J. Vuckovic, and A. Scherer. Design and fabrication of silicon photonic crystal optical waveguides. *J. Lightwave Technol.*, 18:1402–1411, 2000.

[8] M. Tokushima, H. Kosaka, A. Tomita, and H. Yamada. Lightwave propagation through a 120 degrees sharply bent single-line-defect photonic crystal waveguide. *Appl. Phys. Lett.*, 76:952–954, 2000.

[9] K.S. Kunz and R.J Luebbers. *The Finite Difference Time Domain Method for Electromagnetics*. CRC Press, London, UK, 1993.

[10] Berenger J. P. A perfectly matched layer for free-space simulation in finite-difference computer codes. *Ann. Telecomm.*, 51(1-2):39–46, 1996.

[11] S. G. Johnson and J. D. Joannopoulos. Block-iterative frequency-domain methods for maxwell's equations in a planewave basis. *Opt. Express*, 8(3):173–190, 2001.

[12] Lalanne P. Electromagnetic analysis of photonic crystal waveguides operating above the light cone. *IEEE J. Quant. Electron.*, 38(7):800–804, 2002.

[13] P. Bienstman and R. Baets. Optical modelling of photonic crystals and vcsels using eigenmode expansion and perfectly matched layers. *Opt. Quant. Electron.*, 33(4/5):327, 2001.

[14] M. Notomi, A. Shinya, K. Yamada, J. Takahashi, C. Takahashi, and I. Yokohama. Structural tuning of guiding modes of line-defect waveguides of silicon-on-insulator photonic crystal slabs. *IEEE J. Quant. Electron.*, 38(7):736, 2002.

[15] W. T. Lau and S. Fan. Creating large bandwidth line defects by embedding dielectric waveguides into photonic crystal slabs. *Appl. Phys. Lett.*, 81(21):3915–3917, 2002.

[16] T.F. Krauss, B. Vögele, C. R. Stanley, and R. M. De La Rue. Waveguide microcavity based on photonic microstructures. *Ieee Phot. Technol. Lett.*, 9(2):176–178, 1997.

[17] J.S. Foresi, P.R. Villeneuve, J. Ferrera, E.R. Thoen, G. Steinmeyer, S. Fan, J.D. Joannopoulos, L.C. Kimerling, H.I. Smith, and E.P. Ippen. and photonic-bandgap microcavities in optical waveguides. *Nature*, 390:143–145, 1997.

[18] Y. Akahane, T. Asano, B.S. Song, and Noda S. Investigation of high-q channel drop filters using donor-type defects in two-dimensional photonic crystal slabs. *Appl. Phys. Lett.*, 83(8):1512–1514, 2003.

[19] H. Kosaka, T. Kawashima, A. Tomita, M. Notomi, T. Tamamura, T. Sato, and S. Kawakami. Superprism phenomena in photonic crystals. *Phys. Rev. B*, 58:10096–10099, 1998.

[20] H. Kogelnik and C.V. Shank. Coupled-wave theory of distributed feedback lasers. *J. Appl. Phys.*, 43(5):2327–2335, 1972.

[21] M. Meier, A. Mekis, A. Dodabalapu, A. Timko, R.E. Slusher, J. D. Joannopolous, and O. Nalamasu. Laser action from two-dimensional distributed feedback in photonic crystal. *Appl. Phys. Lett.*, 74(1):7–9, 1999.

[22] Susa N. Threshold gain and gain-enhancement due to distributed-feedback in two-dimensional photonic-crystal lasers. *J. Appl. Phys.*, 89(2):815–823, 2001.

[23] Y. Dumeige, I. Sagnes, P. Monnier, P. Vidakovic, C. Meriadec, and A. Levenson. $\chi^{(2)}$ semiconductor photonic crystals. *J. Opt. Soc. Am B*, 19(9):2094–2101, 2002.

[24] C.L. Tang and P.P. Bey. Phase matching in second harmonic generation using artificial periodic structures. *IEEE J. Quant. Electron.*, QE9(1):9–17, 2002.

[25] M. Soljacic, M. Ibanescu, S.G. Johnson, Y. Fink, and J. D. Joannopoulos. Optimal bistable switching in nonlinear photonic crystals. *Phys. Rev. E*, 66(5):055601, 2002.

[26] M. Soljacic, C. Luo, J. D. Joannopoulos, and S. Fan. Nonlinear photonic crystal microdevices for optical integration. *Opt. Lett.*, 28(8):637–639, 2003.

[27] L. H. Spiekman, Y.S. Oei, E.G. Metaal, F.H. Groen, P. Demeester, and M.K. Smit. Ultrasmall waveguide bends: the corner mirrors of the future? *IEE Proc.-Optoelectron*, 142(1):61–65, 1995.

[28] R. Espinola, R. Ahmad, F. Pizzuto, M. Steel, and R. Osgood. A study of high-index-contrast 90 degrees waveguide bend structures. *Opt. Express*, 8:9, 2001.

[29] R. Ahmad, F. Pizzuto, G. Camarda, R. Espinola, H. Rao, and R. Osgood. Ultracompact corner-mirrors and t-branches in silicon-on-insulator. *IEEE Phot. Technol. Lett.*, 14:1, 2002.

[30] B.E. Little, J.S. Foresi, G. Steinmeyer, E.R. Thoen, S.T. Chu, H.A. Haus, E.P. Ippen, L.C. Kimerling, and W. Greene. Ultra-compact si-sio2 microring resonator optical channel dropping filters. *IEEE Phot. Technol. Lett.*, 10(4):8212–8222, 1998.

[31] T. Shoji, T. Tsuchizawa, T. Watanabe, K. Yamada, and H. Morita. Low loss mode size converter from $0.3\mu m$ square si waveguides to singlemode fibres. *Electron. Lett.*, 38(25):1669–1700, 2002.

[32] K. De Mesel, S. Verstuyft, I. Moerman, P. Van Daele, and R. Baets. Oxide-confined laser diodes with an integrated spot-size converter. *IEEE J. Quant. ELectron.*, 8(6):1372–1380, 2003.

[33] M. Galarza, K. De Mesel, S. Verstuyft, D. Fuentes, C. Aramburu, M. Lopez-Arno, I. Moerman, P. Van Daele, and R. Baets. Mode-expanded $1.55\mu m$ inp-ingaasp fabry-perot lasers using arrow waveguides for efficient fiber coupling. *IEEE J. Sel. Top. Quant. Electron.*, 8(6):1389–1398, 2002.

[34] D. Taillaert, W. Bogaerts, P. Bienstman, T.F. Krauss, P. Van Daele, I. Moerman, S Verstuyft, K. De Mesel, and R. Baets. An out-of-plane grating coupler for efficient butt-coupling between compact planar waveguides and single-mode fibers. *IEEE J. Quant. Electron.*, 38(7):949–955, 2002.

[35] D. Taillaert, H. Chong, P. Borel, L. Frandsen, R. De La Rue, and R. Baets. A compact two-dimensional grating coupler used as a polarization splitter. *IEEE Phot. Technol. Lett.*, 15(9):1249–1251, 2003.

[36] H. Benisty, C. Weisbuch, D. Labilloy, M. Rattier, C. J. M. Smith, T. F. Krauss, R.M. De La Rue, R. Houdre, U. Oesterle, C. Jouanin, and D. Cassagne. Optical and confinement properties of two-dimensional photonic crystals. *J. Lightwave Technol.*, 17:2063, 1999.

[37] H. Benisty, D. Labilloy, C. Weisbuch, C. J. M. Smith, T. F. Krauss, D. Cassagne, A. Braud, and C. Jouanin. Radiation losses of waveguide-based two-dimensional photonic crystals: Positive role of the substrate. *Appl. Phys. Lett.*, 76:532–534, 2000.

[38] W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, and D. De Zutter. Out-of-plane scattering in photonic crystal slabs. *IEEE Phot. Technol. Lett.*, 13(6):565–567, 2001.

[39] S.G. Johnson, S. Fan, P.R. Villeneuve, J.D. Joannopoulos, and L.A. Kolodziejski. Guided modes in photonic crystal slabs. *Phys. Rev. B*, 60:5751–5758, 1999.

[40] E. Chow, S.Y. Lin, S.G. Johnson, P.R. Villeneuve, J.D. Joannopoulos, J.R. Wendt, G.A. Vawter, W. Zubrzycki, H. Hou, and A. Alleman. Three-dimensional control of light in a two-dimensional photonic crystal slab. *Nature*, 407:983–986, 2000.

[41] W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, and D. De Zutter. Out-of-plane scattering in 1d photonic crystal slabs. *Opt. Quant Electron.*, 34(1/3):195–203, 2002.

[42] K.K. Lee, D. R. Lim, L. C. Kimerling, J. Shin, and Cerrina F. Fabrication of ultralow-loss si/sio2 waveguides by roughness reduction. *Opt. Lett.*, 26(23):1888, 2001.

[43] S. Suzuki, M. Yanagisawa, Y. Hibino, and K. Oda. High-density integrated planar lightwave circuits using sio2 geo2 waveguides with a high refractive index difference. *J. Lightwave Technol.*, 12:790–796, 1994.

[44] R.J. Deri and E. Kapon. Low-loss iii-v semiconductor optical waveguides. *IEEE J. Quant. Electron.*, 27(3):626–640, 1991.

[45] P.K. Tien. Light waves in thin films and integrated optics. *Appl. Opt.*, 10:2395–2419, 1971.

[46] T.A Germer. Angular dependence and polarization of out-of-plane optical scattering from particulate contamination, subsurface defects, and surface microroughness. *Appl. Opt.*, 36(33):8798, 1997.

[47] A. Talneau, L. Le Gouezigou, and N. Bouadma. Quantitative measurement of low propagation losses at 1.55 μm on planar photonic crystal waveguide. *Opt. Lett.*, 26(16):1259, 2001.

[48] S. J. McNab, N. Moll, and Y. A. Vlasov. Ultra-low loss photonic integrated circuit with membrane-type photonic crystal waveguides. *Opt. Express*, 11(22):2927–2939, Nov. 2003.

[49] 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. Electron.*, 8(4):928–934, 2002.

[50] H. J. Levinson. *Principles of Lithography*. SPIE, Bellingham, Washington, USA, 2001.

[51] W. Bogaerts, V. Wiaux, P. Dumon, D. Taillaert, J. Wouters, S. Beckx, J. Van Campenhout, B. Luyssaert, D Van Thourhout, and R. Baets. Large-scale production techniques for photonic nanostructures. *Proc. SPIE*, 5225:101–112, 2003.

[52] S.Y. Lin, E. Chow, S. G. Johnson, and J. D. Joannopoulous. Demonstration of highly efficient waveguiding in a photonic crystal slab at the $1.5 - \mu m$ wavelength. *Opt. Lett.*, 25(17):1297–1299, 2000.

[53] R.M. Emmons, B.N. Kurdi, and D.G. Hall. Buried-oxide silicon-on-insulator structures i: Optical waveguide characteristic. *J. Quant. Electron.*, 28(1):157–163, 1992.

[54] T. Trupke, J. Zhao, A. Wang, R. Corkish, and M. A. Green. Very efficient light emission from bulk crystalline silicon. *Appl. Phys. Lett.*, 82(18):2996–2998, May 2003.

[55] N.P. Johnson, D.W. McComb, A. Richel, B.M. Treble, and R.M. De la Rue. Synthesis and optical properties of opal and inverse opal photonic crystals. *Synth. Met.*, 116(1-3):469–473, 2001.

[56] E. Yablonovitch, T.J. Gmitter, and K.M. Leung. Photonic band-structure - the face-centered-cubic case employing nonspherical atoms. *Phys. Rev. Lett.*, 67(17):2295–2298, 1991.

[57] J.G. Fleming, S.Y. Lin, I. El-Kady, R. Biswas, and K.M. Ho. All-metallic three-dimensional photonic crystals with a large infrared bandgap. *Nature*, 417:52–55, 2002.

[58] T. Baba, N. Fukaya, and J. Yonekura. Observation of light propagation in photonic crystal optical waveguides with bends. *Electron. Lett.*, 35:654–656, 1999.

[59] C. J. M. Smith. *Waveguide Photonic Microstructures in III-V Semiconductors*. PhD thesis, Faculty of Engineering, Glasgow University, 1998.

[60] A.J. Ward and J.B. Pendry. A program for calculating photonic band structures, green's functions and transmission/reflection coefficients using a non-orthogonal fdtd method. *Comp. Phys. Comm.*, 128(3):590–621, 2000.

[61] W.C. Sailor, F. M. Mueller, and P. R. Villeneuve. Augmented-plane-wave method for photonic band-gap materials. *Phys. Rev. B*, 57:8819–8822, 1998.

[62] Lalanne P. and E. Silberstein. Fourier-modal methods applied to waveguide computational problems. *Opt. Lett.*, 25(15):1092–1094, 2000.

[63] P. Bienstman. *Rigorous and efficient modelling of wavelength scale photonic components*. PhD thesis, Factulty of Engineering, Ghent University, Ghent, Belgium, 2001.

[64] V.N. Astratov, R.M. Stevenson, I.S. Culshaw, D.M. Whittaker, M.S. Skolnick, T.F. Krauss, and R.M. De La Rue. Heavy photon dispersions in photonic crystal waveguides. *Appl. Phys. Lett.*, 77(2):178–180, 2000.

[65] S.G. Johnson, P. R. Villeneuve, S. Fan, and J.D. Joannopoulous. Linear waveguides in photonic-crystal slabs. *Phys. Rev. B*, 62(12):8212–8222, 2000.

[66] M. Notomi, A. Shinya, K. Yamada, J. Takahashi, C. Takahashi, and I. Yokohama. Singlemode transmission within photonic bandgap of width-varied single-line-defect photonic crystal waveguides on soi substrates. *Electron. Lett.*, 37(5):293–295, 2001.

[67] W.T. Lau and S. Fan. Creating large-bandwidth line defects by embedding dielectric waveguides into photonic crystal slabs. *Appl. Phys. Lett.*, 81:3915, 2002.

[68] T. Felici and H.W. Engl. On shape optimization of optical waveguides using inverse problem techniques. *Inverse Prob.*, 17(4):1141–1162, 2001.

[69] A. Chutinan, M. Okano, and S. Noda. Wider bandwidth with high transmission through waveguide bends in two-dimensional photonic crystal slabs. *Appl. Phys. Lett.*, 80(10):1698–1700, 2002.

[70] T. Felici, A. Lavrinenko, D. Gallagher, and T. Davies. New design rules for planer photonic crystal devices obtained using automatic optimisation, leading to record efficiencies. In *Proc. ECOC*, volume 5, page Tu4.1.2, Rimini, Italy, 2003.

[71] K. Kurokawa. *An introduction to the Theory of Microwave Circuits*. Academic Press, New York, USA, 1969.

[72] S. Boscolo, M. Midrio, and T.F. Krauss. Y junctions in photonic crystal channel waveguides: high transmission and impedance matching. *Opt. Lett.*, 27(12):1001–1003, 2002.

[73] R. Wilson, T. J. Karle, I. Moerman, and T. F. Krauss. Efficient photonic crystal y-junctions. *J. Opt. A*, 5:S76–S80, 2003.

[74] H. Chong and R.M. De La Rue. Photonic crystal hexagonal microcovity with partial hole in-filling for add/drop filters. In *Proc. ECOC*, volume 5, page Tu3.1.5, Rimini, Italy, 2003.

[75] B.S. Song, S. Noda, and T. Asano. Photonic devices based on in-plane hetero photonic crystals. *Science*, 300(5625):1537, 2003.

[76] Y. Akahane, T. Asano, B.S. Song, and S. Noda. High-q photonic nanocavity in a two-dimensional photonic crystal. *Nature*, 425(6961):944–647, Oct. 2003.

[77] A. Yariv, Y. Xu, R.K. Lee, and A. Scherer. Coupled-resonator optical waveguide: a proposal and analysis. *Opt. Lett.*, 24(11):711–713, 1999.

[78] T.J. Karle, D.H. Brown, R. Wilson, M. Steer, and T.F. Krauss. Planar photonic crystal coupled cavity waveguides. *IEEE J. Sel. Top. Quant. Electron.*, 8(4):909–918, 2002.

[79] M. Bayindir, B. Temelkuran, and E. Ozbay. Propagation of photons by hopping: A waveguiding mechanism through localized coupled cavities in three-dimensional photonic crystals. *Phys. Rev. B*, 61(18):11855–11858, 2000.

[80] M. Bayindir and E. Ozbay. Heavy photons at coupled-cavity waveguide band edges in a three-dimensional photonic crystal. *Phys. Rev. B*, 62(4):R2247–R2250, 2000.

[81] L.J. Wu, M. Mazilu, T. Karle, and T.F. Krauss. Superprism phenomena in planar photonic crystals. *IEEE J. Quant. Electron.*, 38(7):915–918, 2002.

[82] L.J. Wu, M. Mazilu, and T.F. Krauss. Beam steering in planar-photonic crystals: From superprism to supercollimator. *J. Lightwave Technol.*, 21(2):561–566, 2003.

[83] A.D. Bristow, J.P.R. Wells, W.H. Fan, A.M. Fox, M.S. Skolnick, D.M. Whittaker, A. Tahraoui, T.F. Krauss, and J.S. Roberts. Ultrafast nonlinear response of algaas two-dimensional photonic crystal waveguides. *Appl. Phys. Lett.*, 83(5):851–853, Aug. 2003.

[84] T. Hirono and M. Kohtoku. Optimized offset to eliminate first-order mode excitation at the junction of sstraight and curved multimode waveguides. *IEEE Phot. Technol. Lett.*, 10:7, 1998.

[85] M.K. Chin and T. Ho. Design and modeling of waveguide-coupled single-mode microring resonators. *J. Lightwave Technol.*, 16(8):1433–1446, 1998.

[86] A. Yariv. Critical coupling and its control in optical waveguide-ring resonator systems. *Photon. Technol. Lett.*, 14(4):483–485, Apr. 2002.

[87] J. Backlund, J. Bengtsson, and Larsson A. Waveguide holograms that provide incoupling and novel additional functions. *Proc. ECIO*, pages 339–342, 1999.

[88] M. M. Spühler, B. J. Offrein, G.-L. Bona, R. Germann, I. Massarek, and D. Erni. A very short planar silica spot-size converter using a nonperiodic segmented waveguide. *J. Lightwave Technol.*, 16(9):1680, 1998.

[89] M. Paulus and O. J. F. Martin. Scattering experiments with a diving cylinder. *Opt. Express*, 9:303–311, 2001.

[90] Y. Tanaka, T. Asano, Y. Akahane, B. Song, and S. Noda. Theoretical investigation of a two-dimensional photonic crystal slab with truncated cone air holes. *Applied Physics Letters*, 82(11):1661–1663, 2003.

[91] S.Y. Chou, P. R. Krauss, and P. J. Renstrom. Imprint lithography with 25-nanometer resolution. *Science*, 272(5258):85–87, 1996.

[92] K. Kincade. Extreme-uv sources are on track for next-generation chip manufacturing. *Laser Focus World*, 39(2):55, 2003.

[93] F. Schellenberg. A little light magic. *IEEE Spectrum*, Sept. 2003.

[94] H. Sewell, J. McClay, P. Jenkins, B. Tirri, D. Flagello, and J. Mulkens. 157nm lithography - window of opportunity. *J. Photopol. Sci. Technol.*, 15(4):569–576, 2002.

[95] B. Jajali, S. Yegnanarayanan, T. Yoon, T. Yoshimoto, I. Rendina, and F. Coppinger. Advances in silicon-on-insulator optoelectronics. *IEEE J. Sel. Top. Quant. Electron.*, 4(6):938–947, 1998.

[96] T.T.H. Eng, J.Y.L. Ho, P.W.L. Chan, S.C. Kan, and G.K.L. Wong. Large core ($\sim 60\mu m$) soi multimode waveguides for optical interconnect. *IEEE Phot. Technol. Lett.*, 8(9):1196–1198, 1996.

[97] A.J. Auberton Herve, M. Bruel, B. Aspar, C. Maleville, and H. Moriceau. Smart-cut(r): The basic fabrication unibond(r) soi wafers. *IEICE Tran. Electron.*, E80C(3):358–363, 1997.

[98] I. Abdulhalim. Theory for double beam interference microscopes with coherence effects and verification using the linnik microscope. *J. Mod. Opt.*, 48(2):279–302, 2001.

[99] R. Wuest, P. Strasser, M. Jungo, F. Robin, D. Erni, and H. Jackel. An efficient proximity-effect correction method for electron-beam patterning of photonic-crystal devices. *Microelectron. Eng.*, 67(8):182–188, 2003.

[100] J. Arentoft, T. Sondergaard, M. Kristensen, A. Boltasseva, M. Thorhauge, and L. Frandsen. Low-loss silicon-on-insulator photonic crystal waveguides. *Electron. Lett.*, 38(6):274–275, 2002.

[101] T. Tsuchizawa, T. Watanabe, E. Tamechika, T. Shoji, K. Yamada, J. Takahashi, S. Uchiyama, S. Itabashi, and H. Morita. Fabrication and evaluation of submicron-square si wire waveguides with spot-size converters. In *LEOS Annual Meeting*, page 287, Glasgow, UK, 2002.

[102] J.M. Vaughan. *The Fabry-Perot interferometer*. Adam Hilger, Bristol, 1989.

[103] A. Talneau, M. Mulot, S. Anand, and Ph. Lalanne. Compound cavity measurement of transmission and reflection of a tapered single-line photonic-crystal waveguide. *Appl. Phys. Lett.*, 82(16):2577–2579, 2003.

[104] D. Hofstetter and R. L. Thornton. Measurement of optical cavity properties in semiconducto lasers by fourier analysis of thr emission spectrum. *IEEE J. Quant. Electron.*, 34(10):1914–1923, 1998.

[105] D. Hofstetter and R. L. Thornton. Loss measurements on semiconductor lasers by fourier analysis of the emission spectra. *Appl. Phys. Lett.*, 72(4):404–406, 1998.