Rigoureuze en efficiënte modellering van golflengteschaal fotonische componenten

Rigorous and efficient modelling of wavelength scale photonic components

ir. P. Bienstman

Proefschrift tot het bekomen van de graad van
doctor in de toegepaste wetenschappen: elektrotechniek
Academiejaar 2000-2001

Universiteit Gent
Vakgroep Informatietechnologie

Promotor: prof. Roel Baets

Universiteit Gent
Faculteit Toegepaste Wetenschappen

Vakgroep informatietechnologie
Sint-Pietersnieuwstraat 41
B-9000, Gent, België

Tel: +32 9 264 33 16
Fax: +32 9 264 35 93

Examencommissie:

Prof. D. De Zutter (voorzitter)
Prof. J. Debacker (secretaris)
Prof. R. Baets (promotor)
Prof. P. De Visschere
Prof. B. Dhoedt
Dr. D. Gallagher, Photon Design plc, Oxford, UK
Prof. P. Lagasse
Prof. G. Morthier
Prof. I. Veretennicoff, VUB

Dit werk kwam tot stand in het kader van een mandaat van Aspirant bij het Fonds voor Wetenschappelijk Onderzoek (FWO) - Vlaanderen (België)

This work was carried out in the context of a doctoral fellowship from the Fund for Scientific Research (FWO) - Flanders (Belgium)

Contents

The following is an excerpt from a conversation between two wizards at the Unseen University, discussing the study of the fundamentals of magic.

Finally Cutangle spoke, very slowly and carefully.

'I look at it like this,' he said. 'Before [...], I was like everyone else. You know what I mean? I was confused and uncertain about all the little details of life. But now,' he brightened up, 'while I'm still confused and uncertain it's on a much higher plane, d'you see, and at least I know I'm bewildered about the really fundamental and important facts about the universe.'

Treatle nodded. 'I hadn't looked at it like that,' he said, 'but you're absolutely right. [It has] really pushed back the boundaries of ignorance. There's so much about the universe we don't know.'

They both savoured the strange warm glow of being much more ignorant than ordinary people, who were ignorant of only ordinary things.

"Equal Rites" - Terry Pratchett

Preface

It has been said countless times before, but that does not make it any less true: a PhD thesis like the one you are currently holding would not have possible without the help, contribution and support from many people.

Obviously, I would like to acknowledge the Fund for Scientific Research (FWO) - Flanders, for their support through a doctoral fellowship.

I would also like to thank prof. Lagasse, the head of the INTEC department, for providing the necessary infrastructure and a dynamic context to perform research in.

I want to express my gratitude towards my advisor, prof. Baets, not only for making the photonics group a fascinating and a stimulating environment, but also for giving me the full freedom to explore this research topic.

I specifically want to mention prof. Daniël De Zutter, prof. Frank Olyslager and Henk Derudder from the INTEC electromagnetics group. They pioneered the use of PML in eigenmode expansion and they were always very helpful when discussing the more subtle points of Maxwell's equations.

Obviously, I do not want to forget the legendary gang of building 39. It would be quite dangerous to try and name them all, since I do not want to risk forgetting anyone because of the numerous shifts in occupation of this building. Thanks guys, for the creative coffee breaks, for the Applied Mechanics Lab Sessions on the dynamics of spherical objects in rectangular enclosures, for the infamous worst-joke contests, in short for providing a great atmosphere that makes work very agreeable and that sparks creativity.

I also want to mention Wim and Dirk, and in a later stage also Bert, Frederik, Mathias and Ilse for acting as early testers of my software. I am quite sure you often cursed me when CAMFR spit out another set of frustrating NaNs, but your feedback has been invaluable.

Additionally, Peter Galle, Kristien De Meulder and their predecessors played an important role by keeping our IT infrastructure running smoothly.

I owe a lot to all the people involved in the European COST268

project, not only for their contribution to the different comparison papers, but also for the stimulating discussions during and after the meetings, and for great fun during the extracurricular evening activities.

Bart Demeulenaere also deserves special words of gratitude. During the first months of my PhD he was patient enough to answer literally dozens of my emails on his VCSEL modelling software. It was also the exposure to his work that led me to discover my own interest in scientific computing that eventually gave rise to the birth of CAMFR.

Finally, I would like to thank all my friends and family, especially my mother, who stood by me during these years with their continuing support, kindness and friendship.

Peter Bienstman

Gent, February 8th, 2001.

Nederlandse Samenvatting - Dutch Summary

In de volgende bladzijden wordt een beknopt overzicht gegeven van de belangrijkste verwezenlijkingen in het kader van voorliggend doctoraatswerk. Voor verdere details en uitgebreidere achtergrondinformatie verwijzen we uiteraard naar de Engelse tekst.

1. Inleiding

Achtergrond

De laatste jaren heeft de telecommunicatie-industrie een explosieve groei gekend, denken we maar aan de exponentiële toename van het gebruik van internet en mobiele telefonie. Het is in belangrijke mate de opto-elektronica die een dergelijke evolutie heeft mogelijk gemaakt, aangezien alle verbindingen in het ruggengraatnetwerk optisch verlopen over glasvezelverbindingen. We moeten echter vaststellen dat op middellange termijn de huidige stand van de technologie niet meer zal volstaan om aan de eisen van de telecomsector te voldoen wat betreft bandbreedte en capaciteit. Nieuwe evoluties in de opto-elektronica dringen zich dus op.

Een mogelijke route bestaat erin om het schakelproces in de netwerkknopen, dat vandaag de dag nog elektronisch gebeurt, langs optische weg te laten verlopen. Om dit op een efficiënte manier mogelijk te maken, moeten echter zeer compacte fotonische geïntegreerde circuiten (PIC's) ontworpen worden, die toelaten een verscheidenheid aan optische functionaliteit op een kleine oppervlakte te combineren. Vandaag de dag staat de PIC-technologie nog grotendeels in zijn kinderschoenen. Vooral het aspect compactheid laat nog veel te wensen over, zeker in vergelijking met de integratiedichtheden die gehaald worden in elektronische chips.

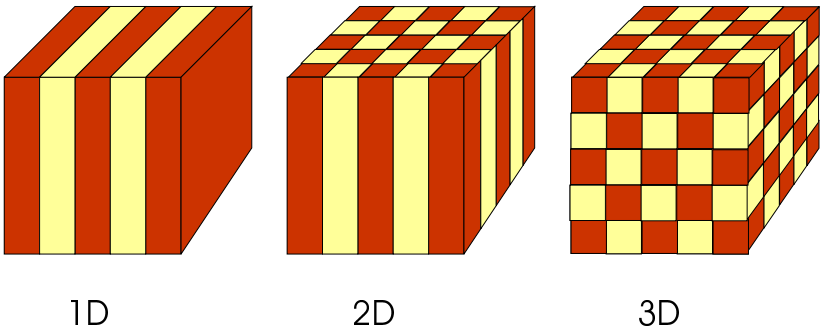

Figuur 1: Fotonische kristallen.

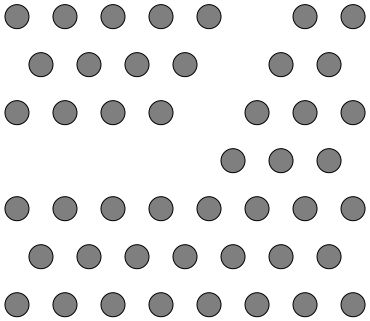

Figuur 2: Bocht in een golfgeleider in een 2D fotonisch kristal.

Een veelbelovende recente ontwikkeling die hier een oplossing zou kunnen bieden is het gebruik van fotonische kristallen (Ph.C.'s) (fig. 1). Dit zijn nanotechnologiecomponenten die een structurele periodiciteit vertonen op een schaal van de golflengte. Door deze periodiciteit zal licht meervoudig verstrooid worden aan de scheidingsvlakken tussen de verschillende materialen. Dit kan aanleiding geven tot destructieve interferentie in de transmissie, zodat een Ph.C. zich als een zeer goede reflector gaat gedragen. Dit principe van een fotonische bandkloof kan aangewend worden om golfgeleiders te maken door opzettelijk de periodiciteit van het kristal te verstoren (door b.v. een rij weg te laten). Op die manier ontstaat een lijndefect in het kristal dat aan beide zijden is omgeven door een perfecte reflector en dus zal fungeren als een golfgeleider. Bovendien kan dit principe uitgebreid worden om in de golfgeleider zeer compacte bochten te nemen (fig. 2), die slechts een oppervlakte van een paar vierkante micron in beslag nemen. Bochten in traditionele golfgeleiders daarentegen zijn typisch vierkante millimeters groot, zodat Ph.C.'s zeer veelbelovend zijn om compacte optische chips te maken. Deze mogelijkheid tot miniaturisatie opent enorme perspectieven, zeker als we kijken welke verbeteringen elektronische

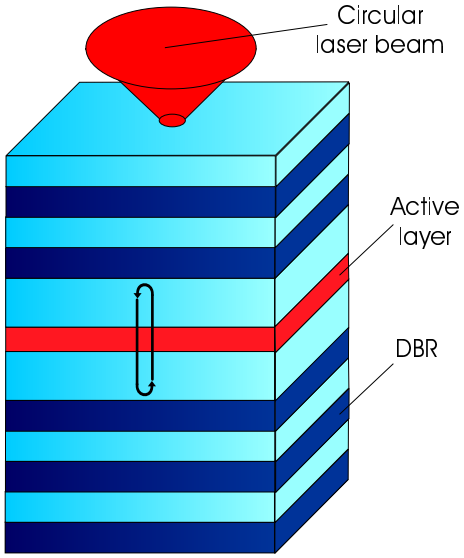

Figuur 3: Verticale-caviteitslaser.

chips ondergaan hebben dankzij een ver doorgedreven miniaturisatie.

Aan de uiteinden van het netwerk, in de laatste kilometer die het ruggengraatnetwerk met de eindgebruiker verbindt, is er ook ruimte voor verbetering. Daar zal de oplossing bestaan uit het doortrekken van optische vezel tot aan de eindgebruiker. Om deze oplossing economisch haalbaar te maken zijn efficiënte en goedkope lasers nodig. Kandidaat bij uitstek hiervoor is de verticale-caviteitslaser (VCSEL) (fig. 3), die nu al een commerciële realiteit is en b.v. zal ingezet worden bij Gigabit Ethernet toepassingen. In een dergelijke laser plant het licht zich voort in de verticale richting, loodrecht op het substraat. Dit staat in tegenstelling tot meer klassieke lasers, waar het licht evenwijdig aan het substraat wordt uitgestraald. Groot voordeel van verticale emissie is dat verschillende lasers op die manier in een tweedimensionale matrix kunnen worden samengebracht op een enkele chip (fig. 4). Bovendien kunnen dergelijke lasers veel compacter gemaakt worden dan hun klassieke tegenhangers, wat dan weer perspectieven opent naar miniaturisatie en snelheid toe.

Modellering van optische componenten

Samengevat kunnen we zonder enige twijfel stellen dat verdere evoluties in de opto-elektronica, en meer bepaald in het domein van PIC's, Ph.C.'s en VCSEL's van cruciaal belang zullen zijn om de verdere ont-

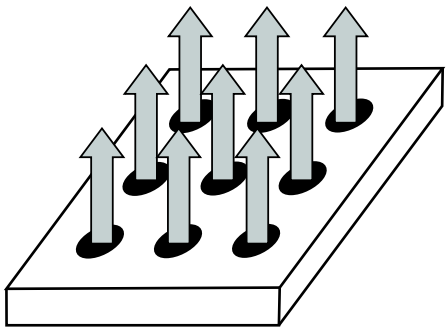

Figuur 4: Optische interconnectie met een 2D matrix van verticaal emitterende lichtbronnen.

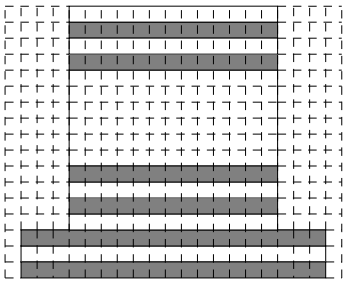

Figuur 5: Ruimtelijke discretisatie.

wikkeling van de telecommunicatie mogelijk te maken. Nochtans is een dergelijke evolutie niet triviaal. In het bijzonder is de fabricage van prototypes bijzonder duur en tijdrovend. Hier dringt zich dus een aanpak op waarbij de verschillende ontwerpsmogelijkheden vooraf doorgerekend worden met behulp van computersimulaties. Probleem hierbij is dat bestaande methodes, zoals b.v. eindige differenties in het tijdsdomein (FDTD), bijzonder traag kunnen zijn en grote hoeveelheden geheugen vereisen. Belangrijkste reden hiervoor is dat deze methodes gebaseerd zijn op het invoeren van ruimtelijke discretisatie (fig. 5). Over de structuur die men wenst door te rekenen wordt een rooster gelegd en men gaat vervolgens op zoek naar de waarde van de elektromagnetische velden ter hoogte van ieder roosterpunt. In de praktijk blijkt men heel veel roosterpunten en dus heel veel onbekenden nodig te hebben. Hierdoor kunnen deze technieken zeer traag zijn en bijgevolg minder bruikbaar voor praktisch ontwerp, waarbij immers verschillende iteraties moeten worden doorgevoerd.

In het kader van dit werk werd dan ook een nieuwe techniek ontwikkeld die precies toelaat om de vereiste rekentijd en hoeveelheid geheugen met een grootteorde te doen dalen. Deze methode is gebaseerd

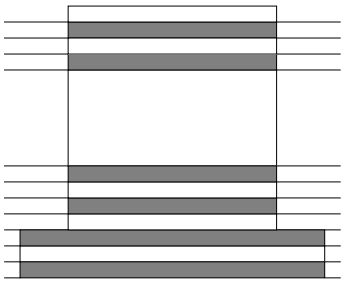

Figuur 6: Eigenmode-expansie.

op een combinatie van vectoriële eigenmode-expansie en het gebruik van geavanceerde randvoorwaarden. De filosofie achter eigenmode-expansie maakt gebruik van het feit dat in het gros van de gevallen de structuur die moet worden doorgerekend een brekingsindexprofiel heeft dat stuksgewijze constant is. Hierdoor kan de component worden verdeeld in lagen waar de brekingsindex niet varieert in de propagatierichting (fig. 6), zodat we in iedere laag het veld kunnen schrijven als een lineaire combinatie van de eigenmodi van die bepaalde laag. Dit resulteert in een veel kleiner aantal onbekenden vergeleken met technieken die de structuur ruimtelijk discretiseren. Hierdoor zijn methodes die gebruik maken van eigenmodi ook merkelijk sneller. Bovendien is eigenmode-expansie rigoureus in de zin dat er geen a priori vereenvoudigingen worden doorgevoerd bij het oplossen van de wetten van Maxwell. Deze aanpak wordt mathematisch rigoureus in de limiet van een oneindig aantal modi en een complete set.

Een probleem dat eigenmode-expansie steeds geplaagd heeft in het verleden en waar in dit werk een oplossing is voor geïntroduceerd, is dat van parasitaire reflecties aan het metalen discretisatievolume. Immers, om een discrete set stralingsmodi te bekomen, geeft men er de voorkeur aan om de structuur die men bestudeert op te sluiten in een metalen doos. Probleem hierbij is dat op die manier de berekeningsresultaten verstoord worden, omdat de straling kan reflecteren aan het metaal en terugkeert naar de structuur die men wil bestuderen. Men kan dit probleem grotendeels oplossen door speciale geavanceerde randvoorwaarden te introduceren, die dan ook in dit werk uitgebreid bestudeerd zullen worden.

Overzicht van het proefschrift

De doelstelling van dit werk is het uitbreiden en verder verfijnen van de eigenmode-expansie om te komen tot een flexibele en performante modelleringstool die toelaat om geavanceerde optische componenten te ontwerpen die de verdere evolutie van de optische telecommunicatie

mogelijk moeten maken.

Dit proefschrift is als volgt opgebouwd. In hoofdstuk 2 wordt een inleiding gegeven op de techniek van eigenmode-expansie. De nieuwe geavanceerde randvoorwaarden die van cruciaal belang zijn voor het succes van deze methode worden vervolgens besproken in hoofdstuk 3. Hoofdstuk 4 gaat dieper in op de software-aspecten gerelateerd aan de implementatie van het model. In de volgende hoofdstukken wordt deze software ingezet bij de studie van een aantal geavanceerde componenten: fotonische kristallen in hoofdstuk 5, verticale-caviteitslasers in hoofdstuk 6 en tenslotte resonante-caviteit-LED's in hoofdstuk 7. Besluiten en mogelijke richtingen voor verder onderzoek worden aangehaald in hoofdstuk 8.

Publicaties

Het werk uitgevoerd in het kader van dit proefschrift heeft geleid tot een aantal publicaties in internationale vaktijdschriften:

- "Detailed study of AlAs-oxidized apertures in VCSEL cavities for optimized modal performance", B. Demeulenaere, P. Bienstman, B. Dhoedt, R. Baets, IEEE Journal of Quantum Electronics, vol. 35, no. 3, pp. 358-367, 1999.
- "Simulation results of transverse-optical confinement in airpost, regrown and oxidized vertical-cavity surface-emitting lasers", P. Bienstman, B. Demeulenaere, B. Dhoedt, R. Baets, Journal of the Optical Society of America B, vol. 16, no. 11, pp. 2055-2059, 1999.
- "The RC^2LED: a novel resonant-cavity LED design using a symmetric resonant cavity in the outcoupling reflector", P. Bienstman, R. Baets, IEEE Journal of Quantum Electronics, vol. 36, no. 6, pp. 669-673, 2000.
- "Analysis of cylindrical waveguide discontinuities using vectorial eigenmodes and perfectly matched layers", IEEE Transactions on Microwave Theory and Techniques, vol. 49, no. 2, pp. 349-354, 2001.
- "Optical modelling of photonic crystals and VCSELs using eigenmode expansion and perfectly matched layers", P. Bienstman, R. Baets, accepted for publication in Optical and Quantum Electronics.
- "Out-of-plane scattering losses in 1D photonic crystal slabs", W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, D. De Zutter, accepted for publication in Optical and Quantum Electronics.

- “Out-of-plane scattering in photonic crystal slabs”, W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, D. De Zutter, accepted for publication in IEEE Photonics Technology Letters.
- "Comparison of optical VCSEL models on the simulation of oxide-confined devices", P. Bienstman, R. Baets, J. Vukusic, A. Larsson, M. Noble, M. Brunner, K. Gulden, P. Debernardi, L. Fratta, G.P. Bava, H. Wenzel, B. Klein, O. Conradi, R. Pregla, S. Riyopoulos, J.-F. P. Seurin, S.L. Chuang, submitted to IEEE Journal of Quantum Electronics.
- "Rigorous and efficient optical VCSEL model based on vectorial eigenmode expansion and perfectly matched layers", P. Bienstman, R. Baets, draft in preparation for IEEE Journal of Quantum Electronics.

Ook is dit werk gepresenteerd op een aantal internationale conferenties:

- "Waveguide and resonator modeling based on vectorial eigenmode expansion and perfectly matched layers", P. Bienstman, R Baets, *invited paper*, Progress In Electromagnetics Research Symposium (PIERS) 2000, p. 175, Cambridge, Massachusetss, USA, July 5-14, 2000.
- "Rigorous optical VCSEL modelling based on vectorial eigenmode expansion", P. Bienstman, B. Demeulenaere, B. Dhoedt, R. Baets, *invited paper*, IEEE LEOS Summer Topical Meetings, paper ThA1.2, San Diego, July 28-30, 1999.
- “Out-of-plane scattering losses in photonic crystal slabs”, W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, *postdeadline paper*, IEEE LEOS annual symposium, PD 1.5, Puerto Rico, November 13-16, 2000.
- “Out-of-plane scattering losses in 1D photonic crystal slabs”, W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, IEEE LEOS Benelux symposium, pp. 35-38, Delft, October 30, 2000.
- "VCSEL modelling based on vectorial eigenmode expansion and perfectly matched layers", P. Bienstman, R. Baets, European workshop on VCSELs, paper M1, Brussels, August 28-30, 2000.
- "Optical VCSEL modelling based on vectorial eigenmode expansion and perfectly matched layers", P. Bienstman, R. Baets, IEEE LEOS Benelux Symposium, pp. 155-158, Mons, November 15, 1999.

- "Microcavity and PBG-related activities at the University of Gent", P. Bienstman, W. Bogaerts, R. Bockstaele, D. Delbeke, S. Goeman, B. Depreter, J. Derluyn, C. Sys, L. Vanwassenhove, B. Dhoedt, I. Moerman, P. Van Daele, R. Baets, COST 268 workshop, Stockholm, March 15-17, 1999.
- "Photonic crystal slabs: how to reduce out-of-plane scattering losses", W. Bogaerts, P. Bienstman. D. Taillaert, R. Baets, RUG-FTW Ph.D. symposium, December 5, 2000, paper 110.
- "Preliminary results on high total dose testing of semiconductor photonic sources: a comparison of VCSELs and resonant cavity LEDs", F. Berghmans, S. Coenen, F. Vos, M.C. Decréton, G. Verschaffelt, A. de la Fuente, I.P. Veretennicoff, H. Thienpont, R. Bockstaele, P. Bienstman, C. Sys, B. Dhoedt, I. Moerman, R.G. Baets, P. van Daele, J. Jonsson, SPIE annual meeting, July 19-24, 1998, San Diego, vol. 3440, pp. 47-56.

Verder hebben we onze medewerking verleend aan een hoofdstuk in een boek over planaire RCLED's:

- "Basics of dipole emission from a planar cavity", R. Baets, P. Bienstman, R. Bockstaele, book chapter in "Confined photon system - fundamentals and applications", pp. 38-79, Springer-Verlag, Heidelberg, 1999.

Het door ons geïntroduceerde concept van de RC^2LED heeft aanleiding gegeven tot de volgende patentaanvraag:

- "A system for guiding a beam of electromagnetic radiation", P. Bienstman, R. Baets, European patent application 99870231.0.

Tenslotte is ook volgend artikel verschenen in de populariserende literatuur:

- "Met de snelheid van het licht" (in Dutch), P. Bienstman, Het Ingenieursblad (KVIV), vol. 68, no. 11-12, pp. 64-69, 1998.

2. Eigenmode-expansie

Eigenmodi in een laag

Deze methode is gebaseerd op het feit dat er in z-invariante media een set speciale velddistributies kan gevonden worden die hun vorm behouden terwijl ze zich langs de z-as voortplanten. Deze zgn. eigenmodi stemmen overeen met de natuurlijke trillingsvormen die kunnen bestaan in dergelijke structuren, ook in de afwezigheid van externe bronnen.

In vele omstandigheden blijkt deze verzameling van eigenmodi een complete set te vormen die kan gebruikt worden om een willekeurig elektromagnetisch veld weer te geven:

$$\begin{cases} \mathbf{E}(\mathbf{r}) &= \sum A_i \mathbf{E}_i(\mathbf{r}_t)\, e^{-j\beta_i z} \\ \mathbf{H}(\mathbf{r}) &= \sum A_i \mathbf{H}_i(\mathbf{r}_t)\, e^{-j\beta_i z} \end{cases} \tag{1}$$

We bekomen op die manier een zeer compacte voorstelling van dit veld als een kolomvector van expansie-coëfficiënten:

$$(\mathbf{E}(\mathbf{r}), \mathbf{H}(\mathbf{r})) \longleftrightarrow \mathbf{A} = [A_i] \tag{2}$$

Verstrooiing aan de scheiding tussen twee lagen

Een volgende stap in het model bestaat erin om uit te rekenen wat er gebeurt met een willekeurige velddistributie $\mathbf{A}_{inc}$ op het ogenblik dat die invalt op het scheidingsoppervlak tussen twee verschillende lagen (fig. 7). Meer bepaald gaan we op zoek naar de reflectie- en transmissiematrices die resp. het teruggekaatste en het doorgelaten veld beschrijven volgens

$$\mathbf{A}_{refl} = \mathbf{R}_{I,II} \cdot \mathbf{A}_{inc} \tag{3}$$

$$\mathbf{A}_{trans} = \mathbf{T}_{I,II} \cdot \mathbf{A}_{inc} \tag{4}$$

Men kan deze matrices bekomen door voor de onbekende velden ook een expansie voor te stellen onder de vorm van eigenmodi. Het toepassen van de gepaste continuïteitsvereisten levert dan uiteindelijk deze matrices, na het in rekening brengen van de orthogonaliteit van de eigenmodi.

Verstrooiing aan een stapel van verschillende lagen

Tenslotte kunnen ook reflectie- en transmissiematrices worden opgesteld die het verstrooiingsgedrag beschrijven van een stapel die bestaat

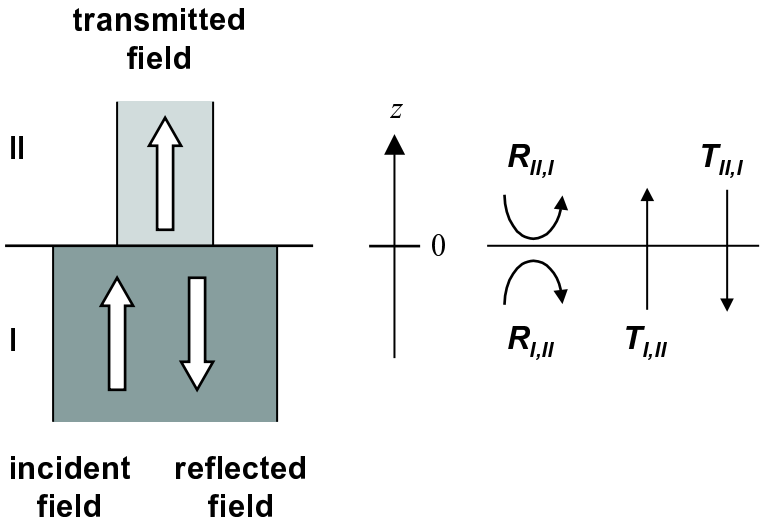

Figuur 7: Het scheidingsoppervlak tussen twee lagen.

uit een opeenvolging van verschillende lagen. Deze matrices worden bekomen door de bijdragen te combineren van de propagatie door de individuele lagen en de verstrooiing die optreedt aan het scheidingsoppervlak tussen twee verschillende lagen. De precieze manier waarop deze bijdragen gecombineerd worden is belangrijk voor de uiteindelijke numerieke stabiliteit van de methode. In de praktijk geven we er daarom meestal de voorkeur aan om te werken met methodes gebaseerd op S-matrices, in plaats van een meer voor de hand liggende beschrijving in termen van transfer T-matrices.

Deze manier om een willekeurige structuur te beschrijven levert ons alle informatie die we nodig hebben. Naast de kennis van het reflectie- en het transmissiegedrag kan ze ook gebruikt worden om het veldprofiel op een willekeurige plaats in de structuur te berekenen.

Zoals reeds aangehaald heeft een dergelijke manier van werken verschillende voordelen ten opzichte van technieken gebaseerd op ruimtelijke discretisatie. Doordat in veel grotere mate gebruik wordt gemaakt van analytisch bekende deeloplossingen (met name in het gedrag van de eigenmodi) kunnen we komen tot een zeer beknopte beschrijving van de velden, een klein aantal onbekenden en een gereduceerde rekentijd. Bovendien laat deze methode in tegenstelling tot andere ook aanzienlijke optimalisaties toe wanneer een basisstructuur optreedt die een aantal maal periodiek wordt herhaald. Klassieke methodes hebben in dit geval een rekentijd die lineair oploopt met het aantal periodes, terwijl voor eigenmode-expansie bepaalde technieken kunnen worden aangewend waardoor de rekentijd slechts logaritmisch toeneemt.

3. Randvoorwaarden

Metalen discretisatiewanden

Realistische structuren zoals die in de praktijk voorkomen zijn open structuren, hetgeen wil zeggen dat straling die wordt uitgezonden door de component naar oneindig kan ontsnappen. Nu blijkt het zo te zijn dat open structuren een continue set stralingsmodi hebben, zodat de velden niet meer beschreven kunnen worden als een discrete som van eigenmodi, en men zijn toevlucht moet nemen tot integralen. Deze integralen zijn niet triviaal om te behandelen, wegens het voorkomen van singulariteiten in het integrand. Vandaar dat men meestal de voorkeur geeft aan een andere aanpak, nl. het plaatsen van de structuur die men wil bestuderen in een gesloten metalen volume.

De aanwezigheid van dit volume discretiseert de set stralingsmodi, maar introduceert terzelfdertijd echter een bijkomend probleem. Immers, de straling die wordt uitgezonden kan nu niet meer naar oneindig propageren, maar zal reflecteren aan het metaal en teruggestuurd worden naar de component. Dit probleem van zgn. parasitaire reflecties kan onder bepaalde omstandigheden de berekeningsresultaten grondig verstoren. Het gebruik van meer geavanceerde randvoorwaarden dringt zich dus op.

Perfect aangepaste lagen

Een zeer interessante geavanceerde randvoorwaarde is het gebruik van perfect aangepaste lagen (PML's), die de eigenschap vertonen straling te kunnen absorberen zonder bijkomende parasitaire reflecties te introduceren. Bovendien is dit het geval voor om het even welke golflengte, invalshoek of polarisatie van de invallende straling, hetgeen een enorm voordeel is ten opzichte van meer traditionele absorberende randvoorwaarden.

In dit werk hebben we gebruik gemaakt van het formalisme van complexe coördinaten om de PML te beschrijven, hetgeen inhoudt dat de PML bekeken wordt als een laag met een reële brekingsindex maar met een complexe dikte. Het is precies het imaginaire deel van deze laagdikte die verantwoordelijk is voor de reflectieloze absorptie in de PML. We kunnen een gevoel krijgen voor de werking van PML door de formule te bekijken die een vlakke golf beschrijft in een uniform medium:

$$e^{-jk_0 nz} \tag{5}$$

Deze golf kan gedempt worden door propagatie in een verlieshebbend medium (n imaginair), maar gezien n en z op gelijke voet optreden in deze formule kan demping ook bekomen worden door propagatie over

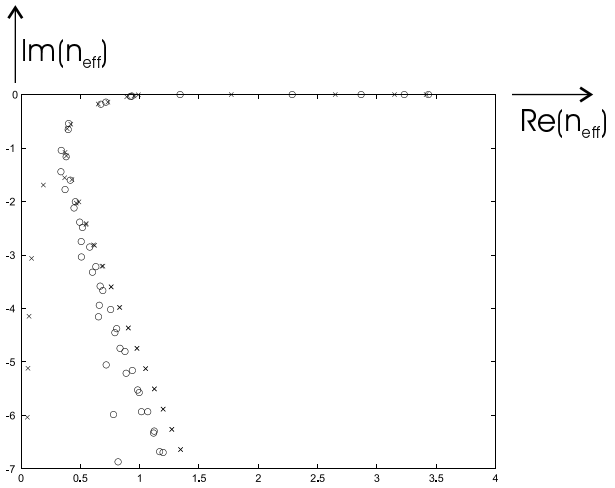

Figuur 8: TE (cirkels) en TM (kruisen) propagatiefactoren van een GaAs ($n = 3.5$) golfgeleider 1 μm dik, omgeven door een luchtlaag met dikte 2.0-0.4j μm. De golflengte is 1.55 μm.

een complexe afstand. Voordeel van een complexe afstand t.o.v. een verlieshebbend materiaal is dat er geen parasitaire reflecties blijken op te treden aan het scheidingsoppervlak tussen de PML en de mantellagen van de bestudeerde component.

Men kan zich terecht vragen stellen bij de fysische betekenis van een dergelijke schijnbaar *ad hoc* geïntroduceerde uitbreiding van de vergelijkingen van Maxwell naar complexe coördinaten. Men kan echter aantonen dat een dergelijke beschrijving equivalent is met een formulering aan de hand van anisotrope media met bijzondere materiaaltensoren, die wel voldoen aan de normale vergelijkingen van Maxwell. Groot praktisch voordeel van de techniek gebaseerd op complexe coördinaten is echter dat alle bestaande analytische formules zeer eenvoudig kunnen worden aangepast om rekening te houden met PML door voor de mantellagen complexe diktes toe te laten.

Fig. 8 illustreert de invloed van PML op de distributie van de propagatieconstantes in het complexe vlak. De geleide modi blijken nog steeds dicht bij assen te liggen, de stralingsmodi bevinden zich nu echter in het complexe vlak ten gevolge van de absorptie in de PML. Deze stralingsmodi kunnen onderverdeeld worden in twee groepen. Een eerste groep (de zgn. Bérenger modi) zal zich steeds verder in het complexe vlak begeven wanneer de absorptie in de PML toeneemt. Een tweede groep convergeert naar een limietwaarde voor toenemende absorptie. Het blijkt dat de locatie van deze limietwaarde precies overeenstemt met de lekkende modi van de corresponderende open structuur. Dit suggereert een dieper verband tussen PML modi en lekkende

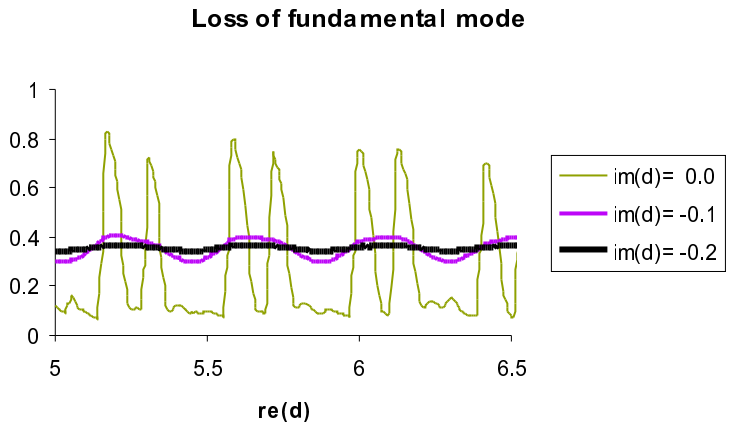

Figuur 9: Stralingsverlies voor de structuur uit fig. 3.3 voor verschillende absorptieniveaus in de PML.

modi. Zo hebben we trouwens kunnen aantonen dat in het geval van een oneindige absorptie PML-mode-expansie equivalent wordt aan lekkende-mode-expansie.

Ter illustratie van het positieve effect van PML op de parasitaire reflecties tonen we in fig. 9 de berekende stralingsverliezen voor de structuur van fig. 3.3. Dit verlies is berekend als een functie van de dikte van de laag tussen de metalen wand en de bestudeerde structuur. In het geval van convergentie verwachten we dat een verdere toename van deze dikte geen veranderingen meer met zich zal meebrengen in de berekende resultaten, aangezien deze dikte een zuiver numerieke parameter is. In de afwezigheid van PML (een imaginaire dikte gelijk aan nul) treden echter sterke oscillaties op in deze curves, zodat het moeilijk wordt om te voorspellen tot welke eindwaarde het verlies zal convergeren. Dit is een duidelijke indicatie van de aanwezigheid van parasitaire reflecties, die aanleiding geven tot resonantie-effecten die de berekening verstoren. Laten we echter de imaginaire component van de dikte oplopen in absolute waarde, dan zien we dat deze resonanties gedempt worden en de storende invloed van parasitaire reflecties verdwijnt.

Geavanceerde randvoorwaarden als PML hebben niet alleen als voordeel dat de nauwkeurigheid gevoelig wordt verbeterd, ze hebben ook een positieve invloed op de rekentijd. Immers, deze absorberende randvoorwaarden laten toe om de randen dichter tegen de bestudeerde structuur te plaatsen, hetgeen resulteert in een kleiner computationeel volume. Dit betekent dan weer dat er minder eigenmodi nodig zullen zijn om de velden te beschrijven. Aangezien de rekentijd oploopt met de derde macht van het aantal modi staat dit garant voor een aanzienlijke snelheidswinst.

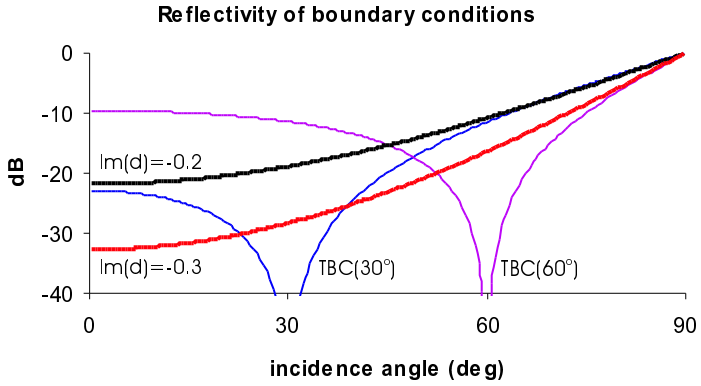

Figuur 10: Vergelijking van verschillende absorberende randvoorwaarden.

Transparante randvoorwaarden

Een tweede klasse van absorberende randvoorwaarden die we bestudeerd hebben wordt gevormd door de zogenaamde transparante randvoorwaarden (TBC's), die een reflectie vertonen die identisch gelijk is aan nul voor een enkele invalshoek. TBC's leggen het volgende verband op tussen voorwaarts en achterwaarts propagerende velden:

$$(k_x + k_{x,0})\, F - (k_x - k_{x,0})\, B = 0 \tag{6}$$

Hier is k_x de transversale component van de golfvector van de invallende straling en $k_{x,0}$ is een vooraf gekozen waarde van k_x waarbij men wenst dat de randvoorwaarde volledig reflectieloos is.

Ter illustratie van deze randvoorwaarden wordt in fig. 10 de absorptie van 2 verschillende TBC's vergeleken met die van twee verschillende PML's. PML's hebben duidelijk de laagste reflecties voor loodrechte inval. TBC's vertonen een reflectie identisch gelijk aan nul voor een bepaalde invalshoek, maar kunnen relatief hoge reflecties vertonen voor andere hoeken. Zowel PML's als TBC's vertonen echter een reflectie die naar 100% convergeert voor scherende inval. Dit betekent dat modi in de buurt van afsnijding geen voordeel halen uit deze randvoorwaarden.

4. CAMFR

Deze eigenmode-expansie met geavanceerde randvoorwaarden is geïmplementeerd in het computerprogramma CAMFR - staande voor CAvity Modelling FRamework. Hierbij hebben we ons laten leiden door object- en component-georiënteerde technieken om te komen tot een flexibel raamwerk dat gemakkelijk is om te gebruiken, te onderhouden en uit te breiden.

Meer in het bijzonder is de software niet opgevat als een enkel monolitisch geheel, maar veeleer als een verzameling van elementaire bouwblokken die op verschillende manieren kunnen worden samengebracht om simulaties uit te voeren, ook op manieren die origineel niet verwacht waren bij het initiële ontwerp. Dit draagt in belangrijke mate bij tot het raamwerk-karakter van de software en zijn flexibiliteit en uitbreidbaarheid.

Ook hebben we gebruik gemaakt van het feit dat een belangrijk deel van de modellen in eigenmode-expansietechnieken geformuleerd kunnen worden zonder een beroep te doen op de precieze details van de onderliggende geometrie. Door deze modellen te implementeren in termen van abstracte golfgeleiders en abstracte verstrooiers kunnen deze zonder meer verder blijven gebruikt worden wanneer er een nieuwe geometrie (b.v. ellipsvormige golfgeleiders) wordt toegevoegd.

Dit alles is mogelijk gemaakt door onze toevlucht te nemen tot een geavanceerde programmeertaal als C++, die voldoende expressief is om dergelijke constructies op te zetten aan de hand van polymorfisme. Voor de gebruikersinterface van CAMFR is gebruik gemaakt van de bestaande scripttaal Python, die o.m. een naadloze integratie met andere programma's toelaat om de berekende resultaten te visualiseren.

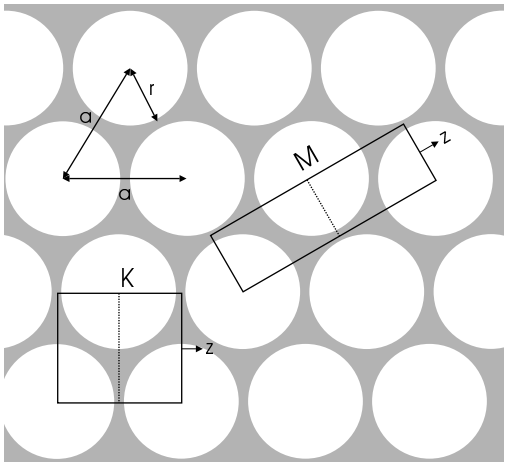

Figuur 11: Triangulair rooster.

5. Fotonische kristallen

Banddiagramma's

Voor de studie van fotonische kristallen is het interessant om te weten welke Blochmodi kunnen bestaan in een bepaalde oneindig uitgestrekte periodieke structuur. Ook hier kunnen we eigenmode-expansie gebruiken, want het blijkt dat we uit de eigenwaarden van de transfermatrix van de basisperiode informatie kunnen halen over de propagatiefactoren van de Blochmodi. De veldprofielen van deze modi volgen dan weer uit de eigenvectoren van deze matrix.

Ter illustratie geven we hier het banddiagramma van een triangulair rooster van luchtgaten in een halfgeleider (fig. 11). Op deze figuur zijn ook de eenheidscellen aangegeven die gebruikt werden bij de berekening van de Blochmodi in een welbepaalde symmetrierichting in de structuur. In de transversale richting van deze eenheidscellen is gebruik gemaakt van elektrische of magnetische wanden om het oneindig uitgestrekt karakter van het kristal weer te geven.

Fig. 12 toont de berekende bandstructuur. Er is duidelijk een frequentieband te zien waar geen modi optreden. Licht dat invalt op het kristal met een frequentie in deze bandkloof zal geen modi vinden om naar te koppelen, zodat het volledig weerkaatst wordt. Hieraan hebben fotonische kristallen hun reputatie als uitstekende spiegels te danken.

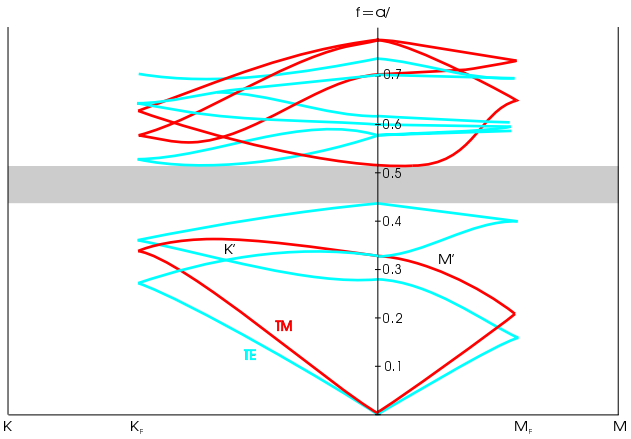

Figuur 12: Bandenstructuur in een triangulair rooster.

Componenten gebaseerd op fotonische kristallen

CAMFR is ingezet bij de studie van een rijke waaier aan componenten gebaseerd op periodieke structuren, gaande van golfgeleiders en bochten in fotonische kristallen, over uitkoppelstructuren en wigmodeconvertoren tot de invloed van brekingsindexcontrast op verstrooiingsverliezen in plaatgolfgeleiders gecombineerd met fotonische kristallen.

Bij wijze van voorbeeld tonen we hier het veldprofiel in een 3dB splitser gecombineerd met scherpe 90-gradenbochten in een fotonisch kristal (fig. 13). Het is duidelijk dat het fotonisch kristal in staat is om het licht te geleiden: het vermogen blijft grotendeels geconcentreerd in de lijndefecten in de kristallen.

We wensen tenslotte ook op te merken dat dankzij eigenmode-expansie dergelijke berekeningen bijzonder snel kunnen gebeuren (2 minuten op een 250 MHz Sun UltraSparc). FDTD methodes zouden hier grootteordes meer rekentijd en geheugen voor nodig hebben. Groot voordeel van eigenmode-expansie in deze context is de mogelijkheid om de structuur onmiddellijk met een Blochmode te exciteren, terwijl FDTD routines moeten vertrekken van een stroomexcitatie die een relatief lang stuk overgangsgolfgeleider vereist om te komen tot een evenwichtsdistributie die enkel de gewenste Blochmode bevat.

Halfoneindige kristallen

In het kader van dit werk is ook onderzoek verricht naar het reflectiegedrag van halfoneindige kristallen. In tegenstelling tot hetgeen courant in de literatuur gesuggereerd wordt, blijkt het in het algemeen niet mogelijk te zijn om te komen tot een eensluitende definitie voor de reflectie

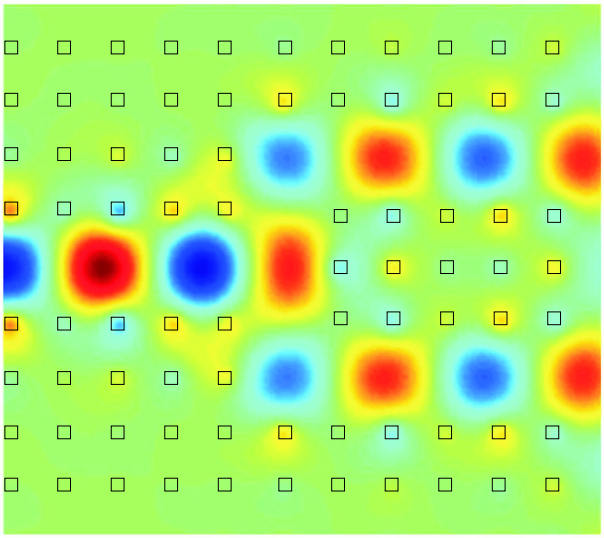

Figuur 13: Halfvermogensplitser in fotonische kristallen.

van dergelijke structuren. In aanwezigheid van geleide Blochmodi blijkt immers deze reflectie te oscilleren als een functie van het aantal periodes, zodat een limietwaarde voor een oneindig aantal periodes niet te definiëren valt.

Dit fenomeen hebben we dan gekoppeld aan de onmogelijkheid om halfoneindige kristallen te beschrijven als een laterale randvoorwaarde, precies omdat dit oscillerende gedrag aanleiding schijnt te geven tot een incomplete set.

Voor meer details verwijzen we uiteraard naar de Engelse tekst.

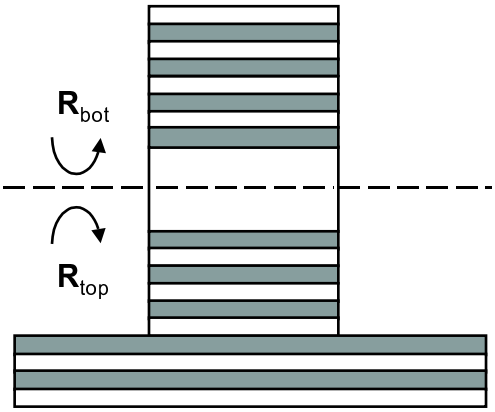

Figuur 14: Lasercaviteit.

6. Verticale-caviteitslasers

Laserresonanties

We kunnen eigenmode-expansie ook gebruiken voor het bepalen van de golflengte en de vereiste materiaalwinst van een lasermode van een willekeurige caviteit. Daarvoor snijden we eerst de caviteit op een arbitraire plaats doormidden (fig. 14). Vervolgens berekenen we de reflectiematrices van de bovenste en de onderste caviteitshelft. Een lasermode moet voldoen aan de conditie dat zijn kringwinst gelijk is aan 1:

$$\mathbf{R}_{top} \cdot \mathbf{R}_{bot} \cdot \mathbf{A}_{lasing} = \mathbf{A}_{lasing} \tag{7}$$

Met andere woorden, we moeten op zoek gaan naar een situatie waarbij de matrix $\mathbf{R}_{top} \cdot \mathbf{R}_{bot}$ een eigenvector heeft met een eigenwaarde gelijk aan 1. Om dit te bekomen variëren we eerst de golflengte λ om aan faseresonatie te voldoen, en vervolgens de materiaalwinst g_{mat} in de actieve laag om amplituderesonantie te bekomen. Op die manier is het vinden van een lasermode herleid tot een zoektocht in het 2D (λ, g_{mat}) vlak.

Transversale opsluiting

Om een laser met een lage drempelstroom te kunnen realiseren is het van groot belang om de laserbundel bij elkaar te houden in een klein gebied in de actieve laag. Als we er dan ook nog eens in slagen om enkel stroom te injecteren in dit gebied, kunnen we een bijzonder efficiënte VCSEL maken.

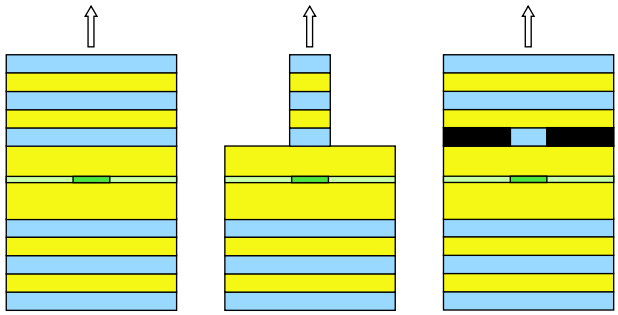

Figuur 15: Proton-geïmplanteerde VCSEL, luchtpaal-VCSEL en VCSEL met dik oxidevenster.

Om deze transversale opsluiting te bekomen is het een groot voordeel om te beschikken over een simulatiemodel als het onze dat het optisch gedrag in dergelijke structuren correct vectorieel doorrekent. Immers, de introductie in de caviteit van een element dat moet zorgen voor transversale opsluiting heeft zowel positieve als negatieve gevolgen. Het is precies deze delicate balans die ervoor zorgt dat het moeilijk is om intuïtief te voorspellen welk effect de opsluitstructuur zal hebben.

Aan de ene kant zal een dergelijke structuur er uiteraard voor zorgen dat de laserbundel bij elkaar wordt gehouden in het *inwendige* van de opsluitstructuur. Dit kan leiden tot een vermindering van de diffractieverliezen, daar waar anders de bundel te veel zou openwaaieren in de caviteit, zodat hij na reflectie aan de spiegels een doorsnede zou hebben die veel groter is dan het gebied waarin stroom geïnjecteerd wordt.

Aan de andere kant kunnen er ook verstrooiingsverliezen optreden aan de *randen* van de opsluitstructuur, wanneer het licht erin gekoppeld moet worden of juist moet verlaten. Dit kan ook een belangrijk verliesmechanisme zijn voor de laserwerking. Vervelend is wel dat structuren die een goede interne opsluiting realiseren vaak ook meer verstrooiingsverliezen veroorzaken door een sterker indexcontrast. Hier moet dus een aanvaardbaar evenwicht gezocht worden.

Ter illustratie vergelijken we nu de drempelwinst van een aantal VCSEL-structuren geschetst in fig. 15. Het gaat om een VCSEL met protonimplantie, een luchtpaal-VCSEL en een VCSEL met een dik ($\lambda/4$) oxidevenster. De proton-geïmplanteerde VCSEL wordt gemaakt door het oppervlak selectief te beschieten met protonen, Hierdoor wordt de kristalstructuur verstoord, zodat een isolator gecreëerd wordt. Op die manier kan een geleidend centraal kanaal gemaakt worden, omgeven door een isolator. Aangezien de protonen geen invloed hebben op de brekingsindex van de lagen, is de enige verandering in brekingsindex het gevolg van de winst in de actieve laag in het gebied onder het geleidend kanaal. Men spreekt in dit verband ook over winstgeleiding.

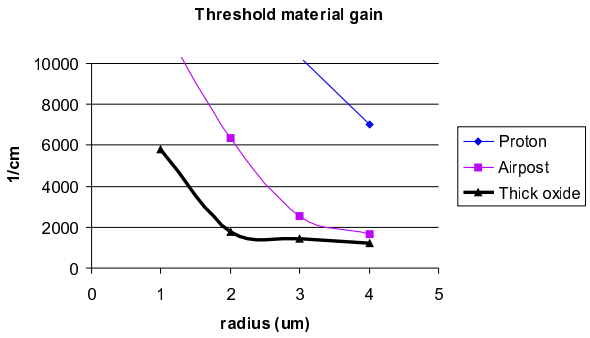

Figuur 16: Drempelwinst van verschillende VCSEL structuren.

Bij de luchtpaal-VCSEL bekomt men transversale opsluiting door de halfgeleiderplak selectief aan te etsen, waardoor een paalvormige structuur blijft staan. Deze fungeert als golfgeleider en kan bijgevolg ook het licht bij elkaar houden. De derde VCSEL-structuur maakt gebruik van de oxidatie van een halfgeleiderlaag met een hoog aluminiumgehalte. Ook dit resulteert in een centraal geleidend kanaal omgeven door een isolerende apertuur van aluminiumoxide. Bovendien heeft dit oxide een lage brekingsindex ($n = 1.55$) zodat ook hier optische golfgeleiding bekomen wordt.

Fig. 16 toont de drempelwinst in deze structuren als een functie van de straal van de component. De winstgeleide VCSEL heeft de hoogste drempelwinst, hetgeen logisch is omdat de transversale opsluiting bijzonder klein is. Zoals reeds gezegd, wordt deze in dit geval enkel bekomen door de winst in de actieve laag. De luchtpaal-VCSEL doet het beter en de oxide-VCSEL heeft de laagste drempelwinst. Nochtans is in de luchtpaal-VCSEL het opsluitingseffect veel groter dan bij de oxide-VCSEL, omdat het brekingsindexcontrast hoger is en het effect zich laat voelen over de hele bovenspiegel in plaats van in een enkele laag. Dit is een duidelijke indicatie van het reeds aangehaalde effect van verstrooiingsverliezen in structuren met een groot brekingsindexcontrast. Hoewel de oxide-VCSEL een lagere opsluiting heeft, zijn ook de verstrooiingsverliezen kleiner wegens het lagere brekingsindexcontrast.

Dunne oxidevensters

Men kan deze redenering verder proberen door te trekken en zich afvragen of men een nog lagere drempelwinst kan bekomen door een dunner oxide te gebruiken. Op die manier offert men wat opsluiting op, maar tegelijkertijd zullen ook de verstrooiingsverliezen afnemen omdat het dunnere oxide de lasermode minder zal verstoren. Bij gebruik

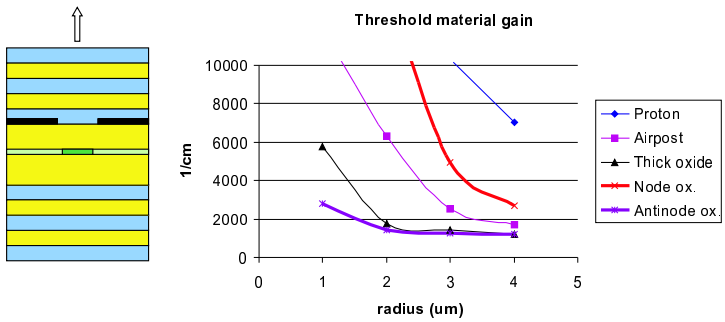

Figuur 17: Drempelwinst van VCSELs met dunne oxidevensters.

van dergelijke dunne oxides krijgt men ook een extra vrijheidsgraad ter beschikking: men kan het oxide positioneren op een minimum (knoop) of een maximum (buik) van het staandegolfpatroon. Fig. 17 toont de simulatieresultaten voor dergelijke dunne oxidevensters. Wanneer het oxide op een maximum van het optische veld wordt geplaatst, blijkt een dergelijke aanpak inderdaad zijn vruchten af te werpen. Het verlies in opsluiting wordt goedgemaakt door de kleinere verstrooiingsverliezen, zodat de component een drempelwinst heeft die lager ligt dan die van een VCSEL met een dik oxidevenster. Voor een dun oxidevenster op een veldminimum liggen de zaken totaal anders. Dit dunne venster heeft bijna geen invloed op het elektromagnetisch veld, precies omdat het zich op een positie bevindt waar dit veld zo zwak is. Dit betekent dat het er niet in zal slagen om de laserbundel bij elkaar te houden, zodat de diffractieverliezen enorm zullen toenemen, hetgeen resulteert in hoge drempelwinsten.

Nochtans mogen dergelijke conclusies niet veralgemeend worden. Zo zijn er immers in de literatuur studies gerapporteerd waarbij precies de tegenovergestelde trend experimenteel werd waargenomen, nl. dat knoopoxides beter presteren dan buikoxides. Dit effect hebben we theoretisch kunnen verklaren door ook de diffusie van ladingsdragers in rekening te brengen. Immers, deze ladingsdragers zullen na injectie in de actieve laag lateraal openspreiden, waardoor de winstzone groter wordt dan de eigenlijke kanaaldiameter. Een dergelijk breder winstprofiel is veel beter aangepast aan de brede lasermode in het geval van knoopoxides, zodat het licht efficiënt gepompt kan worden. Precies het omgekeerde is waar voor buikoxides: veel van de ladingsdragers leveren geen bijdrage tot de winst omdat de lasermode veel smaller is dan de winstoppervlakte. Dit alles wordt schematisch geïllustreerd in fig. 18.

Verder hebben we ook simulaties verricht die illustreren dat het mogelijk is om een dergelijk effect ook positief aan te wenden om een hogere modale stabiliteit te bekomen in VCSELs die een oxidevenster combineren met een protonimplantatie die een kleinere diameter heeft

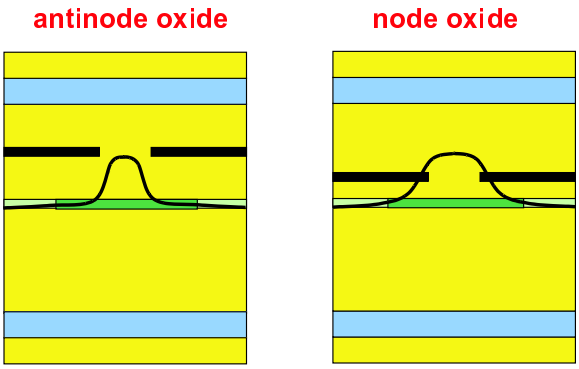

Figuur 18: Invloed van ladingsdragersdiffusie op de overlap tussen het winstprofiel en de lasermode.

dan het oxidevenster. Op die manier is het winstprofiel vooral geconcentreerd in het centrum van de laser, waar de fundamentele lasermode een grote intensiteit vertoont, terwijl die van de hogere orde modes er juist verwaarloosbaar is. Zo wordt de grondmode zeer sterk bevoordeeld, hetgeen resulteert in een betere modale stabiliteit.

Ook is onderzoek gebeurd op componenten die een wigvormig oxidevenster hebben. Doordat in dergelijke lasers het brekingsindexprofiel slechts geleidelijk wijzigt, kunnen de verstrooiingsverliezen nog verder worden teruggedrongen met een nog lagere drempelwinst tot gevolg.

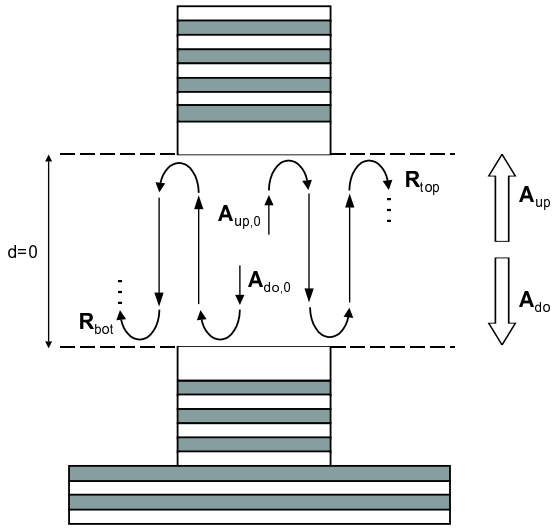

Figuur 19: Bron in een caviteit.

7. Resonante-caviteit-LED's

Respons op een stroombron

Een ander belangrijk probleem is de studie van de respons van een arbitraire caviteit op een willekeurige stroombron. Wanneer we een stroombron gekarakteriseerd door zijn stralingspatroon $(\mathbf{A}_{up,0}, \mathbf{A}_{do,0})$ in een caviteit plaatsen, zal ten gevolge van meervoudige reflecties een gewijzigd stralingspatroon $(\mathbf{A}_{up}, \mathbf{A}_{do})$ ontstaan (fig. 19). Dit patroon wordt gegeven door

$$\mathbf{A}_{up} = (\mathbf{I} - \mathbf{R}_{bot} \cdot \mathbf{R}_{top})^{-1} \cdot (\mathbf{A}_{up,0} + \mathbf{R}_{bot} \cdot \mathbf{A}_{do,0}) \tag{8}$$

$$\mathbf{A}_{do} = (\mathbf{I} - \mathbf{R}_{top} \cdot \mathbf{R}_{bot})^{-1} \cdot (\mathbf{A}_{do,0} + \mathbf{R}_{top} \cdot \mathbf{A}_{up,0}) \tag{9}$$

Een elektrische puntdipool als bronterm is een goed model voor het proces van spontane emissie in een halfgeleider ten gevolge van de recombinatie van een elektron en een gat. We kunnen dus aan de hand van de beschreven vergelijkingen bepalen welke invloed de optische omgeving (i.e. de aanwezigheid van de caviteit) heeft op de spontane emissie. In sommige gevallen kan de caviteit geen modi onderhouden waar de spontane-emissiestraling naar kan koppelen. In dit geval zal de dipool geen licht kunnen uitzenden en spreken we van *onderdrukte spontane emissie*. In andere gevallen bezit de caviteit een mode die precies resonant is met de dipoolstraling. Deze mode kan dan zeer efficiënt

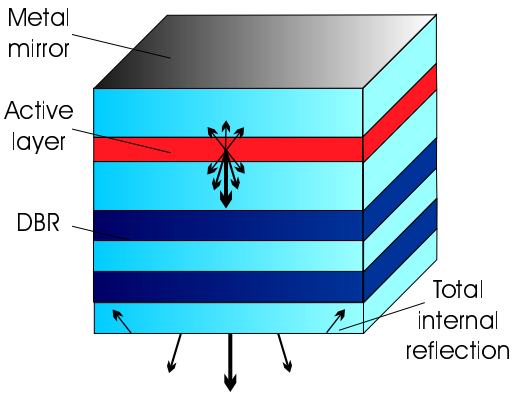

Figuur 20: RCLED.

energie onttrekken aan de dipool zodat we *bevorderde spontane emissie* krijgen.

RCLED

De resonante-caviteit-LED (RCLED) maakt gebruik van dit principe van wijziging van spontane emissie om een grotere extractie-efficiëntie te bekomen (fig. 20). Wanneer straling vanuit het inwendige van een LED onder een te schuine hoek invalt op het scheidingsvlak tussen de halfgeleider en de lucht zal die totaal intern reflecteren. Op die manier kan dit licht niet de buitenwereld bereiken en gaat het dus verloren. Voor klassieke LED's met een isotroop stralingspatroon gaat op die manier 98% van al het licht verloren.

RCLED's maken gebruik van de verandering van spontane emissie om het licht bij voorkeur loodrecht op het substraat uit te zenden. Op die manier speelt totale interne reflectie veel minder een rol en kunnen extractie-efficiënties van boven de 20% gehaald worden, een hele verbetering ten opzichte van de 2% bij klassieke LED's.

Om dit principe van spatiale herverdeling van spontane emissie beter te begrijpen, is het instructief om gebruik te maken van een schematische voorstelling in de k-ruimte (fig. 21). Hier wordt de oorsponkelijk isotrope emissie van de actieve laag voorgesteld door een boloppervlak, terwijl de resonantievoorwaarde van de caviteit weergegeven wordt aan de hand van ander oppervlak. Punten in de k-ruimte waar er veel overlap is tussen beide oppervlakken stemmen overeen met straling die wordt uitgezonden door de caviteit en die door constructieve interferentie wordt versterkt. Uit een dergelijk diagram is dus heel duidelijk te zien dat in een RCLED vooral de straling rond $\theta = 0$ gaat versterkt worden.

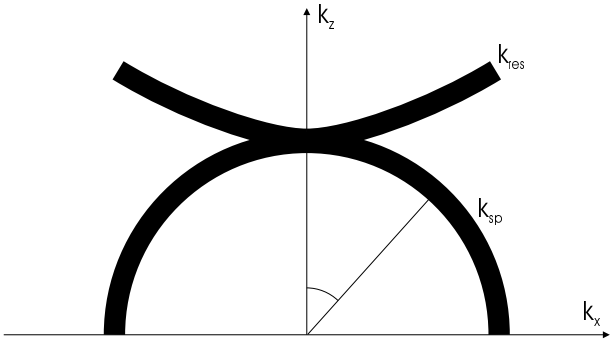

Figuur 21: k-ruimtediagram voor een RCLED.

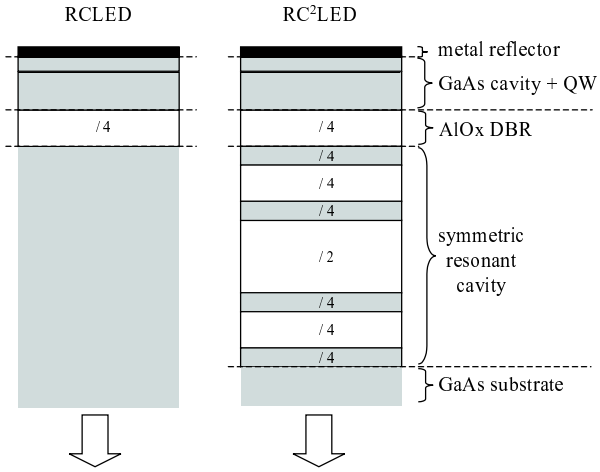

Figuur 22: RC^{2}LED versus RCLED.

RC^{2}LED

In het kader van dit werk hebben we ook een totaal nieuw concept voor dergelijke RCLED's voorgesteld, dat toelaat om een grotere extractie-efficiëntie te halen in combinatie met een smaller stralingspatroon. Deze zogenaamde RC^{2}LED is gebaseerd op het gebruik van een tweede resonante caviteit en is het onderwerp geweest van een octrooiaanvraag. Zoals te zien in fig. 22, is in de RC^{2}LED een extra caviteit geplaatst onder de uitkoppelspiegel.

Deze combinatie van spiegel en caviteit heeft een aantal voordelen. Ten eerste heeft een dergelijke structuur slechts een belangrijke transmissie in een aanzienlijk kleiner hoekinterval. Dit betekent dat het stralingspatroon van deze component veel smaller zal zijn. Bovendien heeft deze spiegel een negatieve angulaire penetratiediepte. Dit

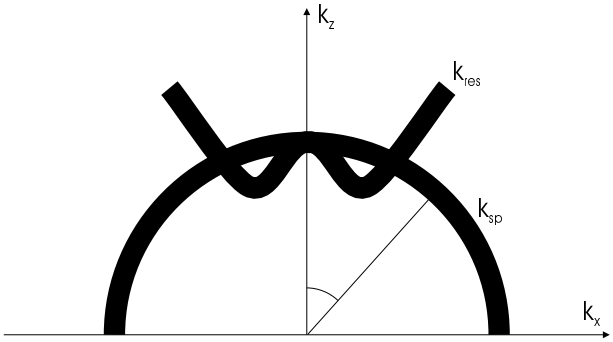

Figuur 23: k-ruimtediagram voor een RC^{2}LED.

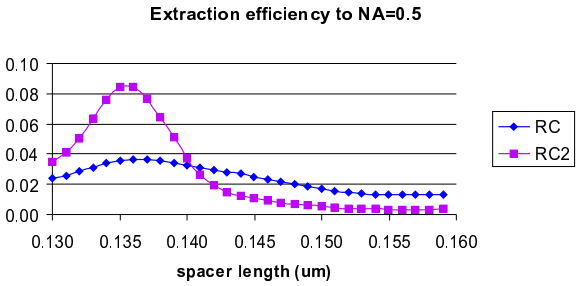

Figuur 24: Extractie-efficiëntie naar NA=0.5 voor een RCLED en een RC^{2}LED.

betekent dan weer dat het resonantieoppervlak in de k-ruimte omgebogen zal zijn (fig. 23), hetgeen resulteert in een grotere totale overlap en een grotere efficiëntie.

Als illustratie van de verbeterde eigenschappen van deze component tonen we in fig. 24 een aantal CAMFR-simulaties waaruit duidelijk blijkt dat de extractie-efficiëntie naar optische vezel minstens een factor 2 hoger kan liggen in vergelijking met een RCLED. Dit betekent dat voor eenzelfde stroom door de component er twee keer zoveel licht in een optische vezel kan terecht komen, hetgeen bijzonder interessant is voor toepassingen als datacommunicatie waarbij het vermogenbudget steeds zeer kritisch is.

8. Besluit

In dit werk hebben we de eigenmode-expansiemethode uitgebreid met geavanceerde randvoorwaarden, zoals perfect aangepaste lagen en transparante randvoorwaarden. Dit laat toe dat de methode op een correcte manier stralingsverlies in rekening kan brengen, daar waar vroeger deze straling volledig reflecteerde aan de harde wanden en op die manier de berekeningsresultaten verstoorde. Dit heeft niet alleen een positieve invloed op de nauwkeurigheid, maar verbetert ook de rekentijd van het model, aangezien de wanden nu veel dichter bij de bestudeerde structuur geplaatst kunnen worden. Dankzij deze reductie in computationeel volume is slechts een kleiner aantal N eigenmodi vereist om tot convergentie te komen, hetgeen een aanzienlijke snelheidswinst met zich meebrengt omdat de rekentijd oploopt als N^3.

We hebben deze technieken geïmplementeerd in een flexibel en gebruikersvriendelijk CGO-raamwerk genaamd CAMFR. Dankzij ver doorgedreven object-georiënteerde technieken kan de implementatie van deze modellen gebeuren zonder expliciete kennis van de onderliggende ruimtelijke geometrie, wat het uiteraard bijzonder gemakkelijk maakt om CAMFR uit te breiden met nieuwe geometrieën.

Op het ogenblik zijn een tweetal van dergelijke geometrieën geïmplementeerd in CAMFR:

- een 2D cartesische geometrie, met een arbitrair aantal indexsprongen in de laterale en de z-richting.
- een 3D cilindersymmetrische geometrie, met een enkele indexsprong in de radiale richting en een arbitrair aantal in de z-richting.

Voor deze geometrieën (en ook voor andere toekomstige geometrieën) is de volgende functionaliteit aanwezig:

- vectoriële veldprofielen in een structuur, inclusief reflectie- en transmissiegedrag voor een willekeurig invalsveld.
- banddiagramma's voor een oneindige periodieke uitbreiding van een basisstructuur.
- lasermodes van een arbitraire caviteit, inclusief hun resonantiegolflengte en drempelwinst.
- respons van een willekeurige caviteit op een willekeurige bron.

Rekentijden in CAMFR vallen bijzonder voordelig uit in vergelijking met andere methodes, vooral dan met deze gebaseerd op ruimtelijke discretisatie zoals FDTD. Dit is grotendeels te danken aan het feit dat

eigenmode-expansie in staat is om gebieden met constante brekingsindex analytisch in rekening te brengen. Dit alles betekent dat rekentijden in CAMFR typisch van de grootteorde seconden of minuten zijn, terwijl FDTD uren of zelfs dagen zou rekenen op dezelfde structuur.

We hebben aangetoond dat CAMFR ingezet kan worden bij het bestuderen van een waaier aan verschillende optische componenten en structuren. We hebben het gebruikt bij de studie van banddiagramma's in fotonische kristallen en om de eigenschappen te berekenen van golfgeleiders, bochten en splitsers in deze materialen. We hebben een gedetailleerde ontwerpsstudie verricht aangaande de verschillende opsluitingsmechanismes in verticale-caviteitslasers en hebben wat meer klaarheid gebracht in de afwegingen die in dergelijke aangelegenheden moeten worden gemaakt. Tenslotte hebben we het proces van spontane emissie in halfgeleiders bestudeerd en hebben we een nieuw concept voor een resonante-caviteit-LED voorgesteld dat een grotere extractie-efficiëntie vertoont naar optische vezel toe.

Wat richtingen voor toekomstig onderzoek betreft, ligt het voor de hand om te denken aan uitbreidingen naar nieuwe geometrieën, zoals circulaire geometrieën met een arbitrair aantal radiale indexsprongen of een volledig 3D cartesische geometrie. Vooral dat laatste is een grote uitdaging, maar wel een veelbelovende onderzoeksrichting, aangezien een dergelijke aanpak potentieel kan leiden tot een bijzonder efficiënte behandeling van 3D optische problemen, iets wat vandaag de dag enkel haalbaar is met behulp van extreme krachtige (en dure) computers.

Chapter 1

Introduction

> "What else is research but a blind date with knowledge?"
> Will Harvey

1.1 Context

These are exciting times. Never before has the rapid evolution of science and technology had a greater impact on our daily lives. This is visible not only in the swiftly advancing medical and life sciences, but especially in the fields of electronics, computing and information technology.

The electronics revolution started off when the transistor was invented just after World War II. For most of its lifetime, the transistor's performance has been doubling every eighteen months, which resulted in the explosive growth of consumer electronics and computers we now take for granted.

Another revolution started a decade later, when Maiman built the first laser. Together with the introduction of the optical fibre in the 1970s, this spurred the rapid take-off of telecommunications. Thanks to this so-called photonics technology, it was possible for the first time to send information over larger distances and with better quality than would have been possible with electronic means.

Very recently, these two revolutions sparked the birth of the internet, which has already made a profound impact on the way we communicate, work or do business and that promises to do so even more in the future. However, the current telecommunications infrastructure will start to feel the strain as more and more people find their way to the internet and start demanding increasing amounts of bandwidth. In fig. 1.1, we plot the evolution of the number of computers connected to the internet, which clearly shows an exponential growth.

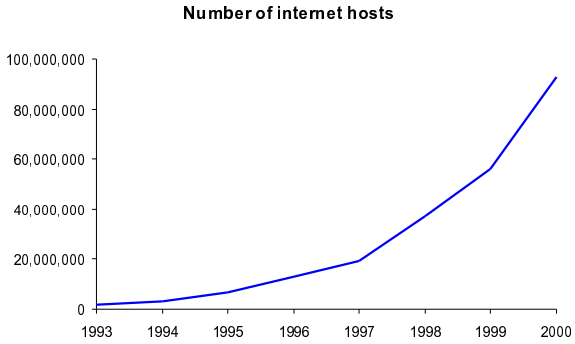

Figure 1.1: Internet host count (source: Internet Software Consortium - www.isc.org).

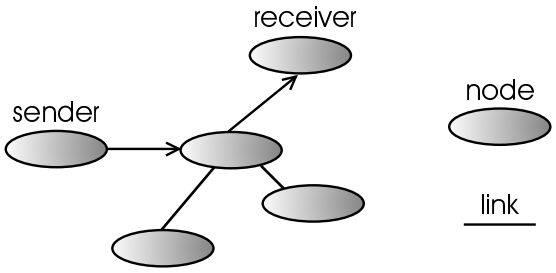

Figure 1.2: Rough sketch of the internet.

The technology underlying the current internet is a mixture of electronics and photonics, a happy marriage of electrons and photons. Fig. 1.2 sketches the rough topology of the internet, and is also representative for a large class of other telecommunications systems.

The information exchanged by computers on the net passes through a number of nodes, which are connected by a series of links. The links transmit the actual information, while the nodes do the routing, i.e. they decide which path the information should follow on its way from source to destination. The links nowadays are mostly optical connections: on one side of the link, data is converted into bursts of light, which are sent through an optical fibre and captured on the other side of the link. While the links are optical, the nodes are still mostly electrical: digital electronics takes the data from the incoming links, converts the light signals to electrical ones, decides to which outgoing links the data has to go and finally launches a new set of optical pulses in the correct outgoing links.

This combination of optical and electronic technology has served the internet well in the past. However, there are more and more indications that the current infrastructure is running out of steam when faced with

the exponential growth of users that we saw in fig. 1.1. This means that there clearly is a need for new technologies, or at least drastic improvements in the old ones.

Currently, the largest bottleneck seems to reside in the electronics. In the past, chip makers have been able to improve the performance of their transistors by miniaturising them. Typical feature sizes on a chip nowadays tend to approach 0.18 μm. However, pushing towards ever smaller feature sizes becomes increasingly harder, not to mention ever more expensive. In the long run, one can already foresee that this miniaturisation cannot continue indefinitely, because eventually transistors will become so small that they will be comparable in size to a single atom, which obviously represents the ultimate limit in miniaturisation.

If we still want to assure a relatively painless growth of the internet over the next few decades, one particularly attractive route would be to give optics a more prominent role in the telecommunications infrastructure. The first place where electronics would have to go is in the nodes, where one could try to do the routing in an entirely optical fashion. Routing photons instead of electrons would eliminate the need for time-consuming conversions from optical signals to electrical signals and back again, which would be a major win.

However, this effectively requires building a photonic chip instead of an electronic one, where the information is carried by photons instead of electrons. This is no mean feat. In the past, considerable progress has already been made in fabricating these PICs (Photonic Integrated Circuits), but they still seem rather primitive and bulky when compared to their electronic counterparts. The main reason why it is so hard to miniaturise these devices, is that it has been very difficult to fabricate structures that allow light to take sharp bends. A typical light bend on a PIC currently takes more space than an entire electronic IC! A technology that would allow for smaller bends would therefore pave the way for the miniaturisation of PICs, and we could expect from that the same kind of drastic improvements that have befallen the electronic ICs in the last decades.

So far, we have focused our attention on the role of optics *inside* the communications network. However, also at the *edges* of the network, in the computers that are connected to it, electronics is running into trouble, especially when it comes to interconnecting different ICs. This is now mostly done electrically, but the main problem with this approach is that two closely spaced interconnections tend to disturb each other, a phenomenon known as cross-talk. This becomes increasingly an issue now that the miniaturisation has driven interconnections closer to each other and now that signals change much faster as the clock rate keeps increasing. Optical interconnections can come to the rescue here, because they are largely immune to these issues. Two light beams can cross or pass arbitrarily close without disturbing each

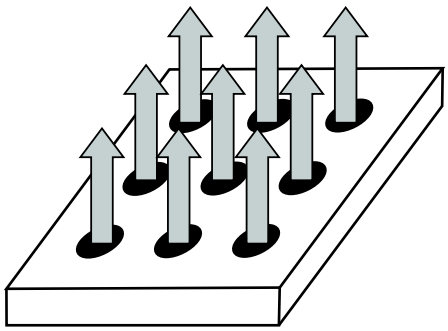

Figure 1.3: Optically interconnected IC.

other, which cannot be said of electrical wires.

Replacing electrical connections between chips by optical ones seems therefore very promising. The largest gain would occur if we could use light sources emitting perpendicularly to the substrate. In this way, we could use the entire surface area of the chip for massively parallel interconnections using a two-dimensional array of light sources (fig. 1.3). These light sources should satisfy a number of constraints. They should be able to be made with a very small diameter, to get a large number of interconnections on a small area. They should also be very efficient, in order to limit power consumption.

In summary, there is a real need for new advances in photonics, which will enable compact PICs and optically interconnected ICs. In the next two sections, we will describe such emerging technologies, which have strong potential to solve a number of outstanding problems facing the telecommunications industry today.

1.2 Photonic crystals

Photonic crystals are the optical analogy of 'regular' crystals. A 'regular' crystal (of a semiconductor material for instance) is a periodic arrangement of atoms or molecules. In such structures, the electronic potential varies periodically in space. Electrons travelling through this lattice will be scattered by this periodic potential. For certain electron energies, the backscattering from the different lattice sites will interfere constructively, so that the electrons cannot penetrate into the crystal. This is the well-known phenomenon of an electronic band gap, which in semiconductors separates the conduction band from the valence band, and has given rise to all sorts of interesting applications, the most prominent of which is the invention of the transistor.

A photonic crystal is based on the same principles, but this time the lattice is formed by a periodic arrangement of macroscopic dielectric

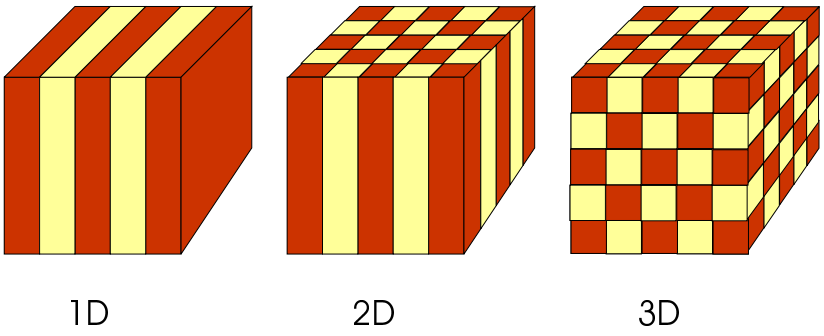

Figure 1.4: Examples of 1D, 2D and 3D photonic crystals.

materials[1], on a length scale several orders of magnitude larger compared to the electronic crystals. The periodicity in space can be either one, two or three-dimensional (fig. 1.4).

One-dimensional photonic crystals are very well-known. In fact, they have been known even before the term 'photonic crystals' was coined, see e.g. [Yeh1977]. A special case of them are none other than the familiar Distributed Bragg Reflectors (DBRs) or quarter-wave stacks, alternating layers of two different materials. When light impinges on such a structure under normal incidence, it will be reflected and transmitted at each of the individual layer interfaces of the stack. By choosing the layer thicknesses equal to a quarter wavelength, all the different contributions to the transmitted beam will interfere destructively, so that most of the light will be reflected backwards. By taking an infinite number of periods, none of the light will be transmitted and we have in fact created a perfect reflector. Or, in electronic crystal parlance, we have created a band gap that forbids the propagation of light under certain circumstances. Because it is a band gap for photons and not for electrons, it is called a Photonic Band Gap (PBG).

Due to their one-dimensional nature, such DBR mirrors only reflect light well for light incident perpendicular to the layers. Light propagating parallel to the layers will not 'feel' any sequential scattering by the layer interfaces and will therefore not be reflected. The obvious solution to this problem is resorting to two-dimensional crystals. When carefully designed, such structures can reflect light coming from all directions, provided that it is propagating in the plane of periodicity of the crystal. Finally, if we want to reflect out-of-plane light too, we will have to use three-dimensional crystals.

Although 1D photonic crystals have been known for a long time, the study of 2D and 3D structures is quite a recent phenomenon. The first

[1] Each of these materials is a 'regular' electronic crystal in its own right, such that a photonic crystal is in a sense a crystal-of-crystals.

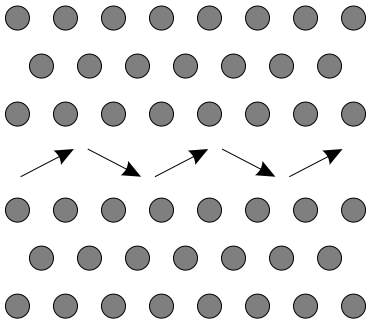

Figure 1.5: Waveguide in a photonic crystal.

suggestion that a material with a 3D PBG might be fabricated appeared in [Yablonovitch1987]. One of the first 3D crystals exhibiting a full photonic band gap was proposed in [Yablonovitch1991]. Fabricating these materials in the 1.5 μm wavelength range suitable for optical communications has only been achieved very recently [Cheng1996]. The nanotechnology skills required for these feats are currently out of reach for most research institutes however.

Photonic band gap materials offer us a way of reflecting light coming from all directions. However, their usefulness does not end there. Just as in electronic band gap materials, they show their full potential when defects that break the perfect periodicity are deliberately introduced in the crystal.

Fig. 1.5 shows the example of a line defect in a 2D photonic crystal, created by omitting a row of rods. Light propagating along the line defect finds itself surrounded on both sides by a photonic band gap material. This PBG forbids the light to enter the crystal, so that it has no other choice than to follow the line defect. We have created in effect a photonic crystal waveguide, which is based on completely different concepts than a traditional optical waveguide. Traditional waveguides are based on the fact that light prefers to travel in regions with high refractive index. In photonic crystals on the other hand, light can be trapped in an air region surrounded by two PBG materials. One obvious advantage of guiding light in air rather than in a high index material, is that air is insusceptible to the non-linear phenomena that plague dielectric materials at high optical powers.

A second advantage of a photonic crystal waveguide is that it allows for light to take a very sharp bend (fig. 1.6). Creating a bend using a conventional waveguide requires a radius of curvature of approximately 1 mm. PBGs are much better at confining the light, such that bends can be fabricated within an area of a few square microns, several orders of magnitude smaller [Mekis1996]. Using this technology, one could therefore start to imagine a miniature optical chip consisting of PBG

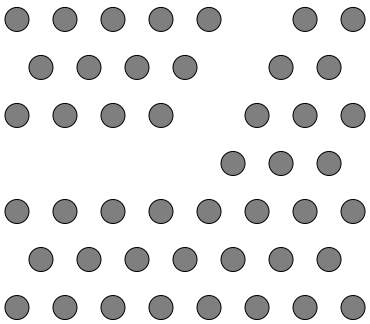

Figure 1.6: Bend in a photonic crystal waveguide.

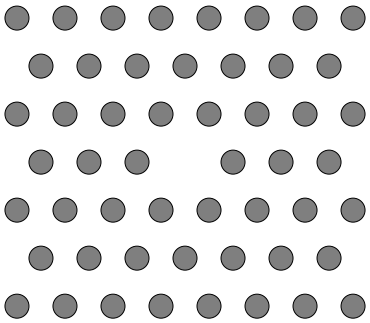

Figure 1.7: Photonic crystal cavity

material, where the light is guided from one place to the next by defects in the crystal lattice.

Going one step further, we can use point defects in PBG materials to create photonic crystal cavities that are completely surrounded by strongly reflecting walls (fig. 1.7). If we could incorporate a light emitter in such a cavity, it would pave the way for very efficient laser structures.

In conclusion, we can say that photonic crystals offer us tremendous possibilities to control the behaviour of light, much more than with conventional technologies. They effectively enable us to 'mold the flow of light', which is also the subtitle of an excellent introductory book on this matter [Joannopoulos1995].

1.3 Microcavity light emitters

Traditional semiconductor lasers have been edge-emitters (fig. 1.8). Inside an active layer in these devices, light is generated which then propagates horizontally, parallel to the substrate. This light bounces back

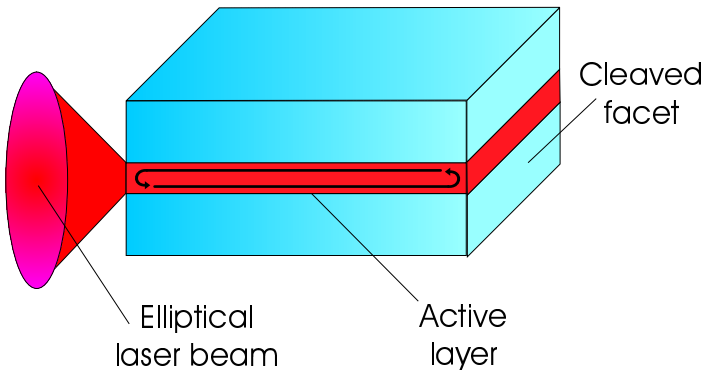

Figure 1.8: Edge-emitting laser.

and forth between two cleaved wafer facets that act as mirrors. These facets typically reflect 30% of the light, and so transmit 70% of the light to the outside world. Not only can the active region emit light, it can also amplify it. This means that if the active region can provide enough amplification, the gain in the cavity can compensate for this loss of light out of the mirrors and laser action ensues. Typical dimensions of such laser diodes are a length of 300 μm and a cross-section of 30 by 30 μm.

These edge-emitters and their cousins have been the workhorses of the optical communications industry for a long time. However, as we already mentioned in section 1.1, it would be very advantageous to have at our disposal a much smaller structure that emits light vertically, perpendicular to the direction of the layers. In the past few years, a device has emerged that does just that: the Vertical Cavity Surface-Emitting Laser or VCSEL (fig. 1.9). A tutorial-like introduction to VCSELs can be found in [Giboney1998], for a more thorough treatment we refer to the book [Sale1995].

In a VCSEL, light propagates vertically instead of horizontally and the role of reflecting facets is now taken over by DBR mirrors, which were already introduced in section 1.2. An important fact to note is that in VCSELs the light only passes through the active layer during a short fraction of its round-trip time. In edge-emitters however, it travels along the active layer during its entire round-trip. For this reason, VCSEL active regions have less opportunity to amplify the optical field, and therefore it is imperative that VCSELs have highly reflecting mirrors to minimise the cavity losses. DBR mirrors used in these devices have typical reflectivities of around 99.9%, which requires a large number of layers.

The VCSEL cavity length is defined by a crystallographic growth step, rather than by a mechanical cleave as was the case with edge-

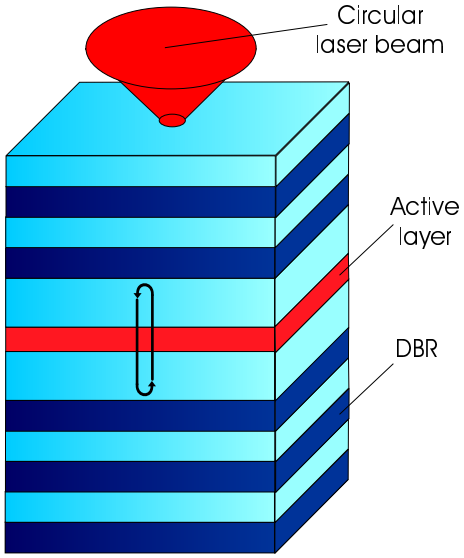

Figure 1.9: Vertical-cavity surface-emitting laser.

emitters. This means that VCSEL cavities can be very short, on the order of the wavelength, which is why they are also called microcavity light-emitters. The main advantage of this short cavity is an inherent longitudinal modal stability, something which requires a complicated design to achieve in edge-emitters.

A further advantage of VCSELs is their wafer-level testability: devices can be tested before the wafer is cut into different chips. This is impossible with edge-emitters, which explicitly require the presence of cleaved facets for their operation.

Finally, because of their surface-emitting character, VCSELs can be integrated into two-dimensional arrays, opening the way for massively parallel optical communication.

VCSELs have improved tremendously over the years, evolving from a laboratory curiosity to a commercial reality. They now achieve performance levels on par or even surpassing those of edge-emitters, making them very interesting for low-cost data communication [Huffaker1999] [Grabherr1999] [Eitel2000].

Another device of the family of microcavity light-emitters is the so-called Resonant-Cavity Light-Emitting Diode or RCLED. It has a structure very similar to that of a VCSEL, but it employs mirrors with a much lower reflectivity (fig. 1.10). One of the mirrors is typically a metal layer, which also doubles as a current injection contact.

In RCLEDs, the mirrors serve an altogether different purpose than

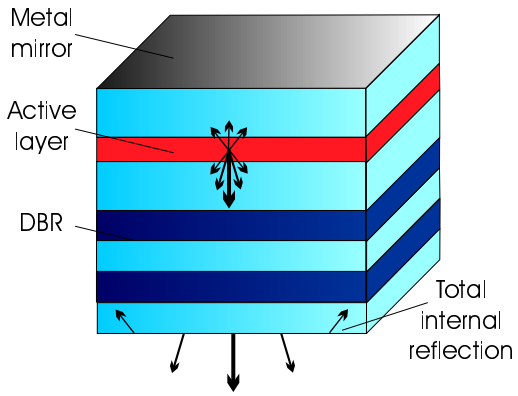

Figure 1.10: Resonant-cavity light-emitting diode.

in VCSELs: their aim is not to provide high reflectivity to achieve lasing, but rather to redistribute the spontaneous light emission in the cavity [Schubert1994]. The multiple reflections of the light off the mirrors cause an interference effect that favours light emission in the direction perpendicular to the mirrors. This is important from a practical point of view, since light that impinges on the mirrors with an angle that is too oblique, will be totally internally reflected and is therefore unable to reach the outside world. In traditional LEDs (i.e. without a micro-cavity) 98% percent of the light emitted is lost in this way. Thanks to their internal redistribution of light, RCLEDs can enjoy a much higher extraction efficiency of up to 20% [De Neve1997c], compared to the 2% in traditional LEDs.

1.4 Optical modelling

Fabricating prototypes of optical components is very time-consuming and costly. This is especially true for today's advanced structures like VCSELs or devices that incorporate photonic crystals. It is therefore no longer feasible to choose the best design by fabricating a large set of possible alternatives and then evaluating them experimentally. This is true in an academic environment, but even more so in a commercial context, where time-to-market constraints demand rapid design cycles. The only possible approach therefore is to resort to computer models that simulate the behaviour of the different designs in an accurate and speedy manner.

In this work we will concentrate on developing models that describe the optical behaviour of complex structures. While the optical characteristics are only part of all the phenomena that govern the device's behaviour (next to e.g. electrical and thermal effects), they neverthe-

less give crucial insight into the component. Moreover, electrical and thermal models are relatively well-developed in literature, while there still is a real need today for optical models that are both accurate and fast at the same time.

There exists a veritable plethora of optical models, and covering them all in depth is beyond the scope of this introduction. For a more detailed overview, we refer to the review articles [Scarmozzino2000] and [Vassallo1997] or the textbook [Itoh1989]. Roughly speaking, optical models can be divided into two different classes. On one hand, there exist a number of models that make certain a-priori assumptions concerning the optical fields. These assumptions allow for a drastic simplification of the equations governing the device behaviour, which leads to models with a very fast run time. However, these assumptions are not always valid and can therefore seriously compromise the accuracy of the obtained results.

On the other hand, there have been a lot of attempts to solve the governing equations exactly, without resorting to any approximation. Due to the general nature of these methods, they can be used to tackle a much larger class of problems, but they typically demand much greater computational resources in terms of memory and running time.

One of the more popular examples in the class of approximative models is the Beam Propagation Method (BPM) [Scarmozzino2000]. This method assumes that the field propagates paraxially, i.e. that it keeps close to a predefined optical axis. Moreover, in its most basic form it neglects the influence from reflections. This means that BPM is mostly restricted to structures with weak refractive index contrast that exhibit low scattering. Using BPM then to model e.g. a bend in a photonic crystal waveguide is therefore out of the question, because PBGs have strong refractive index contrasts and because the bend in the waveguide violates the paraxiality constraint. Nevertheless, the BPM has been widely used to model less demanding structures, like e.g. waveguide tapers or laser facets.

In the class of exact models, most methods employ some form of spatial discretisation, i.e. they overlay the structure to be modelled with a grid and they then solve for the optical field at each of these grid points (fig. 1.11). While these methods are very general, they quickly introduce a large number of unknowns. This typically leads to prohibitively large computational requirements for all but the simplest of structures, sometimes making them almost equally ill-suited to rapid design cycles as prototype fabrication.

Not all methods that use spatial discretisation use the same techniques to solve for the unknowns at the grid points. A first important subclass is that of the finite element methods, that often use a variational technique to minimise an energy functional [Itoh1989]. In the past, these techniques sometimes had the drawback of yielding spurious (i.e. unphysical) solutions, but progress has been made to alleviate

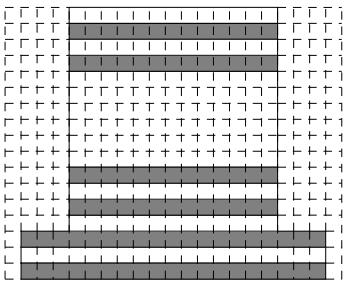

Figure 1.11: Spatial discretisation.

this problem.

A second subclass is based on finite differences, by approximating the derivatives occurring in the formulas by differences. One of the more popular methods in this class is the so-called finite-difference time-domain method (FDTD), which discretises the fields both in the spatial and in the time domain [Taflove1995]. This method too can demand hefty computing resources, but its time domain character has the advantage that a single run with a well-chosen input pulse can yield information on the device's behaviour over a large range of frequencies. Frequency-domain methods can only handle a single frequency at the same time and therefore have to perform multiple runs in order to span a certain frequency range.

Nevertheless, these exact methods are still relatively brute force, so there clearly is a need for another class of models, which also solve the equations exactly, but in far less time. A useful observation in this respect is that the vast majority of the structures to be modelled do not have a totally arbitrary refractive index profile, but rather one where the profile is piecewise constant and does not change over large regions of space. Since the solutions of Maxwell's equations in a uniform space are known analytically, it makes sense to construct the full solution by stitching together the different analytical expressions in each of the individual regions.

Eigenmode expansion techniques do just that. They start out by slicing up the structure into layers where the index profile does not change in a given direction (fig. 1.12).

In each of these individual layers, we can write any optical field as a sum of so-called eigenmodes. An eigenmode of a particular layer can be considered as a natural optical field profile that can exist in this layer and that can propagate indefinitely without changing its form. The expansion of an optical field in terms of eigenmodes is completely analogous to the well-known Fourier expansion in other domains of engineering, where any periodic function can be described as a sum of different sine and cosine functions.

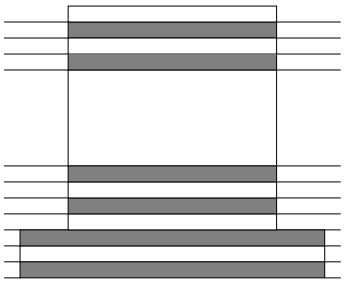

Figure 1.12: Eigenmode expansion.

In this way we create a very compact representation of the optical field, as a small set of weighing coefficients for each of these eigenmodes, rather than as a large number of points where the field is explicitly specified. Eigenmode expansion methods therefore offer the potential of very fast models due to the small number of unknowns.

Despite being fast, these models are still rigorous in the sense that they try to solve Maxwell's equations exactly, without making a priori assumptions like paraxiality or scalar fields. They become mathematically rigorous in the limit of an infinite number of terms in the series expansion using a complete set of eigenmodes.

Eigenmode expansion techniques have already been developed several years ago [Peng1981], [Sudbø1993], [Sztefka1993]. Also at the INTEC department there has been quite a significant amount of work on this topic and related issues, e.g. in the PhD work of [Willems1995], [Haes1996], [Dhoedt1996], and [Demeulenaere1997]. However, eigenmode expansion has so far mainly been applied to small problems, of the kind that can also be modelled with e.g. BPM. At the start of this work, it was still largely an open question whether this method could also efficiently tackle more complicated geometries, like photonic crystal devices or advanced VCSEL designs. More specifically, there was the worrisome matter of the influence of the finite computational domain. Indeed, although optical structures are in reality located in an open infinite space, for numerical purposes we are obliged to truncate this open space, e.g. by placing the structure that we want to study inside a metal box. The problem here is obviously that we are modelling a different structure than the real physical one. This can be an issue because light can be reflected at the metal box and be sent back to the structure and disturb the results.

1.5 Outline of this work

It is the aim of this work to extend and improve the eigenmode expansion method, so that it can be successfully applied to a much larger class of complex optical structures, like photonic crystal devices and advanced VCSEL designs. While still maintaining the same accuracy of the previously mentioned exact models, it should offer substantial improvements in computational efficiency, so that simulation times can be measured in minutes or hours, rather than in days or even weeks using the other methods.

In chapter 2 we will give an introduction to the eigenmode expansion technique. We will classify the different modes and introduce the mode matching method as a way to calculate the scattering matrix that describes the interface between two layers. Different schemes to calculate the scattering matrix of an entire stack of layers will be compared. We will also show how we can calculate the fields at an arbitrary location inside a stack.

Crucial to the success of eigenmode expansion is the use of suitable boundary conditions that will alleviate the problems of parasitic reflections at the boundaries of the computational domain. Therefore, we will give special attention to this topic in chapter 3. Apart from discussing more traditional boundary conditions, like electric or magnetic walls, we will also introduce a recently proposed and very effective boundary condition called a Perfectly Matched Layer (PML). We will discuss the relation of PML to leaky modes. Finally, we will also investigate another kind of walls, namely Transparent Boundary Conditions.

After the more mathematical and physical descriptions in chapters 2 and 3, we will focus in chapter 4 on software engineering aspects related to CAMFR, the eigenmode modelling tool we developed. Special attention will be paid to the concept of abstract waveguides and abstract scatterers, which allow us to rapidly extend the framework with other geometries.

In the next chapters, we will turn to specific application domains for CAMFR and illustrate its usefulness in designing complicated optical structures. Chapter 5 will focus on photonic crystal devices. We will first show how eigenmode expansion can be used to calculate band structures for infinite crystals. Then, we will model some devices like photonic crystal waveguides or bends. We will also investigate whether or not it is possible to introduce special lateral boundary conditions to model semi-infinite crystals. In this context, we will also clear up some common misconceptions appearing in literature about the reflection from semi-infinite photonic crystals.

VCSELs will be the topic of chapter 6. We will discuss some extensions to CAMFR that will allow us to locate the lasing wavelength and the threshold material gain of any cavity. We will use this to perform a thorough design study comparing different forms of transverse optical

confinement. This is crucial when we want to design advanced devices with ultra-low threshold currents.

Chapter 7 will deal with RCLEDs. In order to model these, we will extend CAMFR to calculate the optical response of an arbitrary cavity to an arbitrary current source. In this way, we can model the modification of spontaneous emission rates and extraction efficiencies. We will also discuss a novel design for RCLEDs, which we coined the RC^{2}LED, because it uses two resonant cavities. This design allows us to achieve much higher extraction efficiencies and narrower radiation profiles as compared to traditional designs.

Finally, chapter 8 will offer some conclusions and perspectives for future work.

1.6 Publications

The work carried out during this thesis has led to a number of publications in international refereed journals:

- "Detailed study of AlAs-oxidized apertures in VCSEL cavities for optimized modal performance", B. Demeulenaere, P. Bienstman, B. Dhoedt, R. Baets, IEEE Journal of Quantum Electronics, vol. 35, no. 3, pp. 358-367, 1999.
- "Simulation results of transverse-optical confinement in airpost, regrown and oxidized vertical-cavity surface-emitting lasers", P. Bienstman, B. Demeulenaere, B. Dhoedt, R. Baets, Journal of the Optical Society of America B, vol. 16, no. 11, pp. 2055-2059, 1999.
- "The RC^{2}LED: a novel resonant-cavity LED design using a symmetric resonant cavity in the outcoupling reflector", P. Bienstman, R. Baets, IEEE Journal of Quantum Electronics, vol. 36, no. 6, pp. 669-673, 2000.
- “Analysis of cylindrical waveguide discontinuities using vectorial eigenmodes and perfectly matched layers”, IEEE Transactions on Microwave Theory and Techniques, vol. 49, no. 2, pp. 349-354, 2001.
- “Optical modelling of photonic crystals and VCSELs using eigenmode expansion and perfectly matched layers”, P. Bienstman, R. Baets, accepted for publication in Optical and Quantum Electronics.
- “Out-of-plane scattering losses in 1D photonic crystal slabs”, W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, D. De Zutter, accepted for publication in Optical and Quantum Electronics.

- “Out-of-plane scattering in photonic crystal slabs”, W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, D. De Zutter, accepted for publication in IEEE Photonics Technology Letters.
- "Comparison of optical VCSEL models on the simulation of oxide-confined devices", P. Bienstman, R. Baets, J. Vukusic, A. Larsson, M. Noble, M. Brunner, K. Gulden, P. Debernardi, L. Fratta, G.P. Bava, H. Wenzel, B. Klein, O. Conradi, R. Pregla, S. Riyopoulos, J.-F. P. Seurin, S.L. Chuang, submitted to IEEE Journal of Quantum Electronics.
- "Rigorous and efficient optical VCSEL model based on vectorial eigenmode expansion and perfectly matched layers", P. Bienstman, R. Baets, draft in preparation for IEEE Journal of Quantum Electronics.

This work was also presented at a number of (mostly international) conferences:

- "Waveguide and resonator modeling based on vectorial eigenmode expansion and perfectly matched layers", P. Bienstman, R Baets, *invited paper*, Progress In Electromagnetics Research Symposium (PIERS) 2000, p. 175, Cambridge, Massachusetss, USA, July 5-14, 2000.
- "Rigorous optical VCSEL modelling based on vectorial eigenmode expansion", P. Bienstman, B. Demeulenaere, B. Dhoedt, R. Baets, *invited paper*, IEEE LEOS Summer Topical Meetings, paper ThA1.2, San Diego, July 28-30, 1999.
- “Out-of-plane scattering losses in photonic crystal slabs”, W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, *postdeadline paper*, IEEE LEOS annual symposium, PD 1.5, Puerto Rico, November 13-16, 2000.
- “Out-of-plane scattering losses in 1D photonic crystal slabs”, W. Bogaerts, P. Bienstman, D. Taillaert, R. Baets, IEEE LEOS Benelux symposium, pp. 35-38, Delft, October 30, 2000.
- "VCSEL modelling based on vectorial eigenmode expansion and perfectly matched layers", P. Bienstman, R. Baets, European workshop on VCSELs, paper M1, Brussels, August 28-30, 2000.
- "Optical VCSEL modelling based on vectorial eigenmode expansion and perfectly matched layers", P. Bienstman, R. Baets, IEEE LEOS Benelux Symposium, pp. 155-158, Mons, November 15, 1999.

- "Microcavity and PBG-related activities at the University of Gent", P. Bienstman, W. Bogaerts, R. Bockstaele, D. Delbeke, S. Goeman, B. Depreter, J. Derluyn, C. Sys, L. Vanwassenhove, B. Dhoedt, I. Moerman, P. Van Daele, R. Baets, COST 268 workshop, Stockholm, March 15-17, 1999.
- "Photonic crystal slabs: how to reduce out-of-plane scattering losses", W. Bogaerts, P. Bienstman. D. Taillaert, R. Baets, RUG-FTW Ph.D. symposium, December 5, 2000, paper 110.
- "Preliminary results on high total dose testing of semiconductor photonic sources: a comparison of VCSELs and resonant cavity LEDs", F. Berghmans, S. Coenen, F. Vos, M.C. Decréton, G. Verschaffelt, A. de la Fuente, I.P. Veretennicoff, H. Thienpont, R. Bockstaele, P. Bienstman, C. Sys, B. Dhoedt, I. Moerman, R.G. Baets, P. van Daele, J. Jonsson, SPIE annual meeting, July 19-24, 1998, San Diego, vol. 3440, pp. 47-56.

We also contributed to a book chapter on planar RCLEDs:

- "Basics of dipole emission from a planar cavity", R. Baets, P. Bienstman, R. Bockstaele, book chapter in "Confined photon system - fundamentals and applications", pp. 38-79, Springer-Verlag, Heidelberg, 1999.

The work on the RC^{2}LED was the subject of a patent application:

- "A system for guiding a beam of electromagnetic radiation", P. Bienstman, R. Baets, European patent application 99870231.0.

An introductory paper describing RCLEDs was also published in the periodical of the Royal Flemish Engineer's association:

- "Met de snelheid van het licht" (in Dutch), P. Bienstman, Het Ingenieursblad (KVIV), vol. 68, no. 11-12, pp. 64-69, 1998.

Chapter 2

Eigenmode expansion

> "There is nothing so practical as a good theory."
> Kurt Lewin

2.1 Introduction

Our aim in this chapter is to outline the theory of eigenmode expansion, so that we will be able to calculate the reflection and transmission at an arbitrary stack of arbitrary layers, for an arbitrary incident field.

We will do this in three steps. First, in section 2.2, we will deal with individual layers and describe the properties of the eigenmodes that can exist in these layers. Then, in section 2.3, we will describe how to calculate the scattering that occurs at an interface between two different layers, using the well-known mode-matching method. Subsequently, in section 2.4, we will compare two methods to calculate the reflection and transmission at a stack consisting of an arbitrary number of layers. Calculating field profiles inside a stack will be the subject of section 2.5.

Scattering properties of a layered stack are in many cases not the desired end result of the simulations, but merely a required intermediate step. However, in this chapter we will restrict ourselves to this basic building block and leave more advanced models for chapters 5, 6 and 7.

We will try to describe the scattering at a stack as much as possible without any reference to the underlying layer geometry, i.e. whether the layers are best described in a cartesian coordinate system, a circular one, or indeed any other one. While such a general development of the eigenmode expansion method might seem overly abstract at first, such an approach will prove fruitful when we will consider implementation aspects of the model in chapter 4. In this chapter, we will only include geometry-dependent descriptions in appendices A and B, for slab waveguides and circular waveguides respectively. For a more tutorial-

like introduction to basic waveguide theory, we refer the reader to any standard textbook, like e.g. the excellent [Lee1986].

2.2 Eigenmodes in z-invariant media

2.2.1 Maxwell's equations

Maxwell's equations describe the behaviour of electromagnetic fields. Assuming monochromatic fields that vary harmonically in time according to $e^{j\omega t}$, these equations look like (see e.g. [Chew1990]):

$$\begin{aligned} \nabla \times \mathbf{E} &= -j\omega\mathbf{B} && (2.1)\\ \nabla \times \mathbf{H} &= \mathbf{J} + j\omega\mathbf{D} && (2.2)\\ \nabla \cdot \mathbf{D} &= \rho && (2.3)\\ \nabla \cdot \mathbf{B} &= 0 && (2.4) \end{aligned}$$

By taking the divergence of the curl equations 2.1 and 2.2 and after using the continuity equation $\nabla \cdot \mathbf{J} + j\omega\rho = 0$, we recover the time derivative of the divergence equations 2.3 and 2.4. This shows that the divergence equations can be considered as the initial conditions for the curl equations. Since we restrict ourselves to time-harmonic fields, we do not require initial conditions and can therefore suffice by solving just the curl equations.

In this work, we will consider only linear media, where the following constitutive relations hold:

$$\begin{aligned} \mathbf{D} &= \varepsilon\mathbf{E} && (2.5)\\ \mathbf{B} &= \mu\mathbf{H} && (2.6) \end{aligned}$$

In its most general form, ε and μ are tensors, but we will restrict ourselves to isotropic media where these quantities are scalars.

2.2.2 z-invariant media

Consider the structure in fig. 2.1, where ε and μ do not change in the z-direction and where no sources are present. The volume considered can be bounded by a cylindrical surface δS, oriented along z, but it can also be completely open when δS tends to infinity.

In these structures, there exist field profiles that have the following dependence on z (see e.g. [Vassallo1985]):

$$\begin{cases} \mathbf{E}(\mathbf{r}) = \mathbf{E}(\mathbf{r}_t)\, e^{-j\beta z} \\ \mathbf{H}(\mathbf{r}) = \mathbf{H}(\mathbf{r}_t)\, e^{-j\beta z} \end{cases} \qquad (2.7)$$

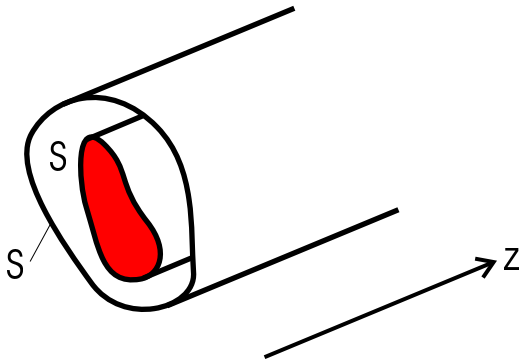

Figure 2.1: z-invariant structure.

Here, $\mathbf{r}$ is the spatial coordinate vector, which can be split up in its longitudinal component z and its transverse component $\mathbf{r}_t$. Equations 2.7 describe the so-called *eigenmodes* of the structure, which are fields that can propagate in the absence of any sources and that maintain their general shape during propagation. β is called the *propagation constant* of the eigenmode. When this quantity is real, the only z-dependent field variation is a phase factor, while the field amplitude remains constant. For complex values of β, the amplitude increases or decreases exponentially, but its shape is maintained. A related quantity is the *effective index* n_{eff} of the mode, defined as

$$\beta = \frac{2\pi}{\lambda} n_{eff}$$

with λ the wavelength of the light.

These field profiles are called eigenmodes, because they satisfy the following mathematical eigenvalue problem, which can be derived from Maxwell's equations:

$$\left(\nabla_t^2 + \omega^2 \mu \varepsilon\right) \mathbf{E} = \beta^2 \mathbf{E} \tag{2.8}$$

Here, β^2 is the eigenvalue. A similar equation can be derived for $\mathbf{H}$.

When we have a set of eigenmodes $\{\mathbf{E}_i, \mathbf{H}_i\}$ at our disposal, we can try to write any forward-propagating field as a linear combination of the modes in this set:

$$\left\{ \begin{array}{lcl} \mathbf{E}(\mathbf{r}) & = & \sum A_i \mathbf{E}_i(\mathbf{r}_t) e^{-j\beta_i z} \\ \mathbf{H}(\mathbf{r}) & = & \sum A_i \mathbf{H}_i(\mathbf{r}_t) e^{-j\beta_i z} \end{array} \right. \tag{2.9}$$

Providing our set of eigenmodes is complete, we can represent any forward-propagating field in this z-invariant structure as an algebraic vector of expansion coefficients:

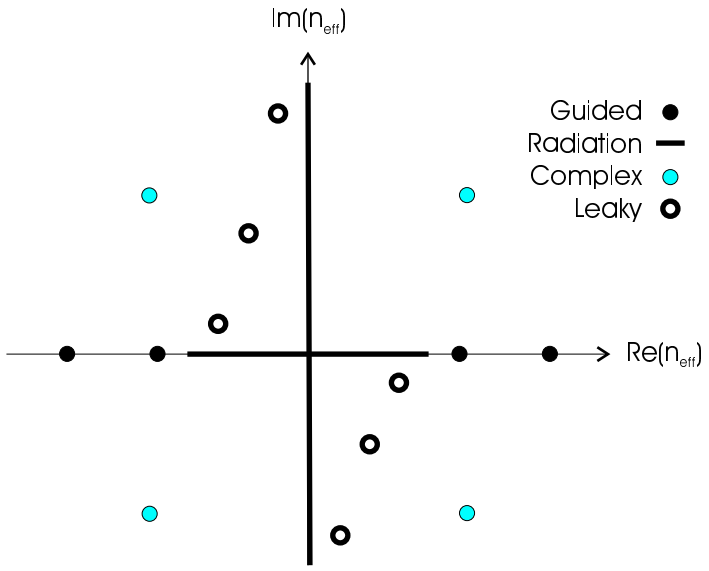

Figure 2.2: Distribution of eigenmodes of an open structure in the n_{eff}-plane.

$$(\mathbf{E}(\mathbf{r}), \mathbf{H}(\mathbf{r})) \longleftrightarrow \mathbf{A} = [A_i]$$

How exactly these eigenmodes are found is not relevant at the moment. This depends on the geometry of the cross-section S and as already mentioned, these aspects will not be discussed until appendices A and B. We will now however briefly touch upon the kind of modes that can exist in general structures like the one in fig. 2.1. We will have to distinguish between open and closed geometries, depending on the presence of a lateral cylindrical boundary δS.

2.2.3 Eigenmodes in open structures

For the sake of argument, we will consider a lossless reciprocal structure with a central core region with refractive index n_{co} and an outer cladding region with a lower index n_{cl}. When plotting the distribution of the different modes in the complex n_{eff}-plane, we get a picture like the one in fig. 2.2.

A first feature of fig. 2.2 is the inversion symmetry around the origin: if β is a propagation constant of a mode, then so is $-\beta$. A propagation constant with a positive real part corresponds to a forward propagating wave, while one with a negative real part is a backward propagating wave with otherwise a similar field profile. This symmetry is not present

in the most general case of non-reciprocal media, but for the situations we will consider it is always true.

Otherwise, the modes can clearly be separated into four groups:

- *Guided modes* are located on the real axis. They form a discrete set, and have effective indices between the cladding index n_{cl} and the core index n_{co}. Their field is concentrated mainly in the high index core of the structure and it decays exponentially in the cladding regions. Because of their real effective index, they propagate undamped in the z-direction. Hence the name guided modes, since they can guide power along the core in the propagation direction.

- *Radiation modes* form a continuum along the real and the imaginary axes. Their field profiles oscillate in the cladding and extend all the way to infinity. This is why they are called radiation modes, because they can carry energy away from the core in the transverse direction. Radiation modes can be *propagative* (real β), in which case their amplitude does not decrease in the z-direction, or *evanescent* (imaginary β), in which case they damp out exponentially along the propagation direction.

- *Leaky modes* are located in the complex plane. They can be seen as a continuation of the guided modes, in the sense that they will become part of the guided modes as the frequency increases. They have a field profile that increases exponentially in the cladding, and are therefore not physical. Nevertheless, they are still useful sometimes to describe fields under certain circumstances [Lee1995] and will be the focus of section 3.4.

- *Complex modes* occur in quartets $\beta = \pm\beta_{re} \pm j\beta_{im}$ in the complex plane. Their occurrence is very rare, and they are only observed when there is a resonant transfer of energy between TE and TM polarised fields [Oliner1981]. It is impossible to excite a single complex mode of a quartet. They are always excited in pairs, in such a way that their combination carries neither active nor reactive power [Omar1987].

The field profiles of the different kinds of modes are sketched in fig. 2.3. The field of guided modes is concentrated in the core, radiation modes show oscillatory behaviour in the cladding. Leaky modes increase exponentially in the cladding, and so do half of the complex modes. The other half of the complex modes (located in the first and third quadrant) decrease exponentially in the cladding, just like the guided modes.

We will now deal with the important aspects of physicality and completeness. A mode is said to be *physical* when the presence of this individual mode can be measured experimentally. Guided modes are

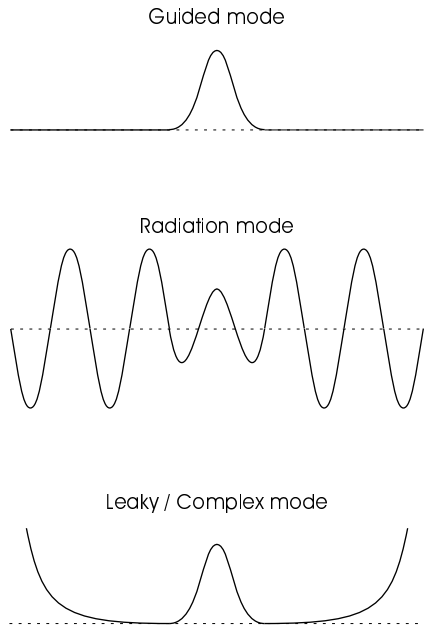

Figure 2.3: Mode profiles of the different kinds of modes.

undoubtedly physical and so is a pair of corresponding complex modes. As already mentioned, a single complex mode is not physical, because it violates power conservation in a lossless medium. Radiation modes or leaky modes[1] are also not physical, because their fields extend all the way towards infinity.

Another aspect, completely separate from physicality, is that of completeness. A set of modes is said to be *complete* when it can describe any field using eigenmode expansion. For structures that are lossless and that cannot support complex modes (like e.g. slab waveguides) it can be rigorously proven[Sagan1989] that the guided modes together with the radiation modes form a complete set. These proofs are based upon the properties of general hermitian operators, like e.g. eq. 2.8 for lossless structures. For all other situations, like e.g. lossy media or waveguides having complex modes, formal completeness proofs currently lack. However, guided and radiation modes are still a set that for all practical intents and purposes appears to be complete, and they are therefore used routinely, despite the absence of rigorous justification. For lossy media, the combination of guided and radiation modes still performs well. When complex modes are present, they have to be

[1] Some measured waveguide phenomena can be described by a single leaky mode, but only in a limited region of space, close to the core. As such, leaky modes are not truly physical in the sense that we described here.

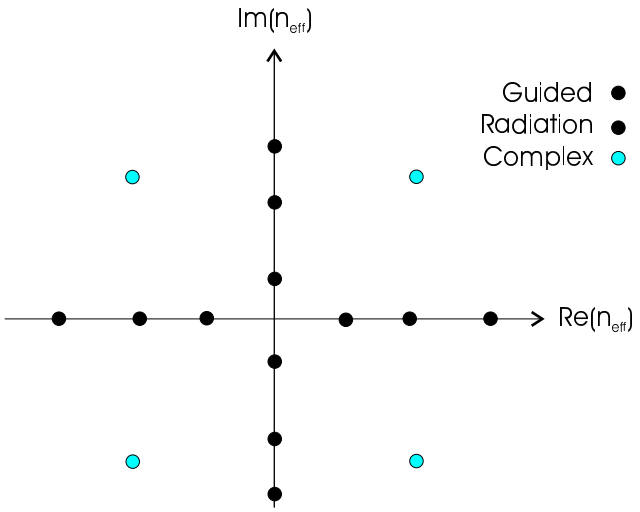

Figure 2.4: Distribution of eigenmodes of a closed structure in the n_{eff}-plane.

added to the set, because otherwise the eigenmode expansion does not converge to the measured results [Chen1988]. Finally, as far as leaky modes are concerned, they are known not to form a complete set, even in combination with the guided modes [Snyder1983].

As said before, the radiation modes form a continuum, meaning that in the simple series expansion of eq. 2.9 the sums have to be replaced by integrals. Subsequently, we can no longer uphold the representation of the field by a discrete set of expansion coefficients A_i. Although it is possible to treat this continuum and the integrals it entails numerically, such an approach is far from trivial due to singularities in the integrand [Chew1990]. Because of this, we will not use this continuum in this work, but rather resort to closed structures, where all modes are discrete, as we will see next.

2.2.4 Eigenmodes in closed structures

For closed structures, the distribution of eigenmodes looks like fig. 2.4.

The presence of the boundary δS imposes additional constraints on the fields, like e.g. the requirement that the tangential electric fields should be zero at the boundary in the case of a perfectly conducting metal wall. These extra conditions make that we no longer have a continuum of radiation modes, but rather a discrete set. In effect, it even becomes questionable to talk about radiation modes, since these modes are in fact guided by the boundary. Nevertheless, it is conventional to refer to the modes with an effective index between n_{cl} and n_{co} as the

guided modes, and the other modes as the radiation modes.

Leaky modes are not present in closed structures. They are a typical open structure phenomenon, with their field profiles increasing exponentially towards infinity. Complex modes however can still occur in closed structures, but they remain relatively rare.

The field profiles of modes in closed structures are qualitatively very similar to the ones in fig. 2.3, but are of course limited to the region inside the boundary.

As far as the completeness is concerned, similar conclusions apply as for open geometries. In case of lossless structures, guided modes together with radiation modes form a complete set.

2.3 Scattering at an interface between two layers

Before tackling the problem of calculating the scattering at an arbitrary interface, we will first derive a very powerful theorem known as Lorentz reciprocity. We will subsequently use this theorem to derive orthogonality relations for eigenmodes.

2.3.1 Reciprocity

Consider two different solutions to Maxwell's equations in the same structure: one solution $(\mathbf{E}_1, \mathbf{H}_1)$, caused by the source $\mathbf{J}_1$, and a second solution $(\mathbf{E}_2, \mathbf{H}_2)$, due to a source $\mathbf{J}_2$. We now write Maxwell's equations for these fields:

$$\nabla \times \mathbf{E}_1 = -j\omega\mu\mathbf{H}_1 \tag{2.10}$$
$$\nabla \times \mathbf{H}_1 = \mathbf{J}_1 + j\omega\varepsilon\mathbf{E}_1 \tag{2.11}$$
$$\nabla \times \mathbf{E}_2 = -j\omega\mu\mathbf{H}_2 \tag{2.12}$$
$$\nabla \times \mathbf{H}_2 = \mathbf{J}_2 + j\omega\varepsilon\mathbf{E}_2 \tag{2.13}$$

Forming the expression $\mathbf{H}_2\cdot(2.10) - \mathbf{E}_1\cdot(2.13) + \mathbf{E}_2\cdot(2.11) - \mathbf{H}_1\cdot(2.12)$, and making use of the vector identity $\nabla\cdot(\mathbf{A}\times\mathbf{B}) = \mathbf{B}\cdot\nabla\times\mathbf{A} - \mathbf{A}\cdot\nabla\times\mathbf{B}$, we finally arrive at

$$\nabla\cdot(\mathbf{E}_1\times\mathbf{H}_2 - \mathbf{E}_2\times\mathbf{H}_1) = \mathbf{J}_1\cdot\mathbf{E}_2 - \mathbf{J}_2\cdot\mathbf{E}_1 \tag{2.14}$$

Integrating both sides over a volume V and making use of Gauss' theorem yields

$$\int\int_S (\mathbf{E}_1\times\mathbf{H}_2 - \mathbf{E}_2\times\mathbf{H}_1)\cdot d\mathbf{S} = \int\int\int_V (\mathbf{J}_1\cdot\mathbf{E}_2 - \mathbf{J}_2\cdot\mathbf{E}_1)dV \tag{2.15}$$

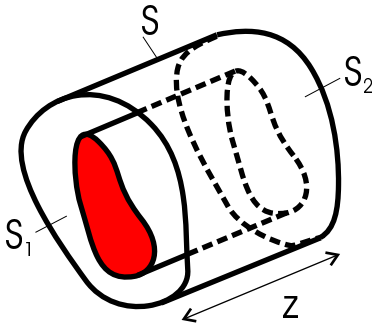

Figure 2.5: Surface over which the Lorentz reciprocity theorem is to be evaluated.

Here, S is the surface that surrounds the volume V. Equation 2.15 is known as the Lorentz reciprocity theorem. For our z-invariant media, we can further simplify this formula by choosing for S the surface indicated in fig. 2.5, a cylinder parallel to the z-axis, with an infinitesimal length Δz.

The surface integral on the LHS of eq. 2.15 can be split into three parts: the contribution from the boundary δS, and from the bottom and top parts S_1 and S_2 of the cylinder. The δS-contribution is determined by the boundary conditions, and those will be extensively treated in chapter 3. Let us suffice at this moment by saying that this term will be zero for most of the cases considered in this work. For guided modes in open structures, this is because of the exponential decay of the fields at infinity. In closed structures with electric or magnetic walls, the boundary contribution vanishes because either the tangential electric or magnetic field is zero at the wall.

Defining

$$\mathbf{F} \equiv \mathbf{E}_1 \times \mathbf{H}_2 - \mathbf{E}_2 \times \mathbf{H}_1 \tag{2.16}$$

we can write for the two remaining terms on the LHS of eq. 2.15

$$\begin{aligned} \int\int_{S_1} \mathbf{F} \cdot d\mathbf{S}_1 + \int\int_{S_2} \mathbf{F} \cdot d\mathbf{S}_2 &= -\int\int_{S_1} F_z dS + \int\int_{S_2} F_z dS \\ &= \int\int_{S} (F_{z+\Delta z} - F_z) dS \end{aligned} \tag{2.17}$$

For the RHS of eq. 2.15, the fields contained in the volume V are independent of z as Δz approaches zero, so the volume integral reduces to a surface integral times the length Δz:

$$\int\int\int_V (\mathbf{J}_1 \cdot \mathbf{E}_2 - \mathbf{J}_2 \cdot \mathbf{E}_1) dV = \Delta z \int\int_S (\mathbf{J}_1 \cdot \mathbf{E}_2 - \mathbf{J}_2 \cdot \mathbf{E}_1) dS \tag{2.18}$$

Equating eq. 2.17 and 2.18, dividing both sides of the equation by Δz and taking the limit $\Delta z \to 0$ then yields the Lorentz reciprocity theorem for z-invariant media:

$$\int\int_S \frac{\partial}{\partial z}\left(\mathbf{E}_1 \times \mathbf{H}_2 - \mathbf{E}_2 \times \mathbf{H}_1\right) \cdot \mathbf{u_z} dS = \int\int_S (\mathbf{J}_1 \cdot \mathbf{E}_2 - \mathbf{J}_2 \cdot \mathbf{E}_1) dS \quad (2.19)$$

2.3.2 Orthogonality

We start from the Lorentz reciprocity theorem for z-invariant media, and take for $(\mathbf{E}_1, \mathbf{H}_1)$ and $(\mathbf{E}_2, \mathbf{H}_2)$ the field distributions of two eigenmodes m and n of the stucture:

$$\begin{array}{l} \left\{ \begin{array}{rcl} \mathbf{E}_m(\mathbf{r}) & = & \mathbf{E}_m(\mathbf{r}_t)\, e^{-j\beta_m z} \\ \mathbf{H}_m(\mathbf{r}) & = & \mathbf{H}_m(\mathbf{r}_t)\, e^{-j\beta_m z} \end{array} \right. \\ \left\{ \begin{array}{rcl} \mathbf{E}_n(\mathbf{r}) & = & \mathbf{E}_n(\mathbf{r}_t)\, e^{-j\beta_n z} \\ \mathbf{H}_n(\mathbf{r}) & = & \mathbf{H}_n(\mathbf{r}_t)\, e^{-j\beta_n z} \end{array} \right. \end{array} \quad (2.20)$$

Eigenmodes can exist in sourceless media, so the RHS of eq. 2.19 vanishes:

$$-j\left(\beta_m + \beta_n\right) \int\int_S \left(\mathbf{E}_m \times \mathbf{H}_n - \mathbf{E}_n \times \mathbf{H}_m\right) \cdot \mathbf{u}_z dS = 0 \quad (2.21)$$

Realising that only the tangential field components contribute, we can write:

$$\left(\beta_m + \beta_n\right) \int\int_S \left(\mathbf{E}_{m,t} \times \mathbf{H}_{n,t} - \mathbf{E}_{n,t} \times \mathbf{H}_{m,t}\right) \cdot \mathbf{u}_z dS = 0 \quad (2.22)$$

We now investigate what happens if we write a similar expression, but this time for the backward propagating mode n. It is relatively easy to show [Lee1986] by splitting Maxwell's curl equations into their transverse and z-components, that for any eigenmode solution

$$\left(\mathbf{E}_{n,t}, \mathbf{E}_{n,z}, \mathbf{H}_{n,t}, \mathbf{H}_{n,z}, \beta_n\right) \quad (2.23)$$

there exists also a second solution corresponding to a backward propagating eigenmode:

$$\left(\mathbf{E}_{n,t}, -\mathbf{E}_{n,z}, -\mathbf{H}_{n,t}, \mathbf{H}_{n,z}, -\beta_n\right) \quad (2.24)$$

Making the necessary substitutions in eq. 2.22 we get:

$$\left(\beta_m - \beta_n\right) \int\int_S \left(\mathbf{E}_{m,t} \times \mathbf{H}_{n,t} + \mathbf{E}_{n,t} \times \mathbf{H}_{m,t}\right) \cdot \mathbf{u}_z dS = 0 \quad (2.25)$$

If we assume that the modes are non-degenerate, such that $\beta_m \neq \beta_n$, we can add equations 2.22 and 2.25 to arrive at the following orthogonality relation:

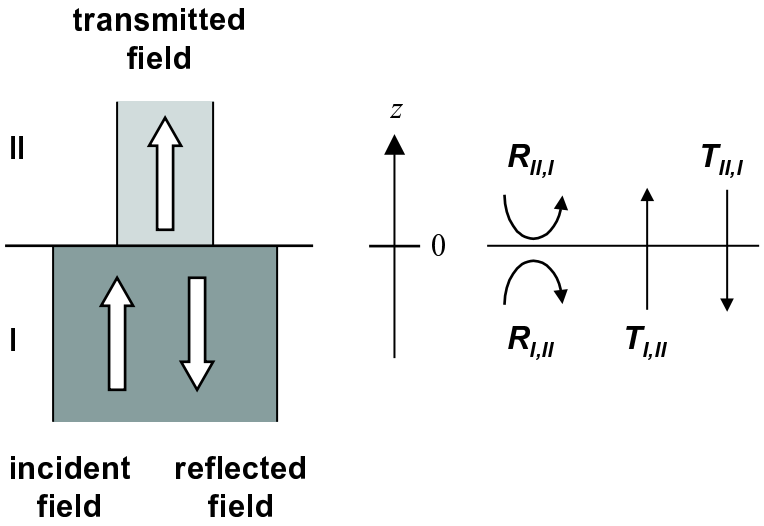

Figure 2.6: Interface between two layers.

$$\int\int_S (\mathbf{E}_{m,t} \times \mathbf{H}_{n,t}) \cdot \mathbf{u}_z dS = 0 \tag{2.26}$$

This will allow us to simplify the expressions when calculating the scattering matrix of an interface, as we will do next.

2.3.3 Mode matching

Consider the flat constant interface between two waveguides I and II, as shown in fig. 2.6. It is placed at $z = 0$ and a single mode with index p is incident from medium I. This incident mode will give rise to a backward-propagating field in medium I, which we expand in terms of the eigenmodes of this medium. Likewise, we expand the transmitted field in the eigenmodes of medium II. The following derivation is none other than the well-known mode-matching technique (see e.g [Zaki1988]). It starts off by imposing the continuity of the tangential components of the total field:

$$\mathbf{E}^I_{p,t} + \sum_j R_{j,p}\mathbf{E}^I_{j,t} = \sum_j T_{j,p}\mathbf{E}^{II}_{j,t} \tag{2.27}$$

$$\mathbf{H}^I_{p,t} - \sum_j R_{j,p}\mathbf{H}^I_{j,t} = \sum_j T_{j,p}\mathbf{H}^{II}_{j,t} \tag{2.28}$$

The minus sign for the reflected $\mathbf{H}$ field is due to the symmetries for the backward-propagating modes already discussed in eq. 2.24. To

calculate the unknown expansion coefficients $R_{j,p}$ and $T_{j,p}$ (which can be seen as reflection and transmission coefficients), we take the right cross product of eq. 2.27 with $\mathbf{H}^I_{i,t}$ and the left cross product of eq. 2.28 with $\mathbf{E}^I_{i,t}$. Here, i is an arbitrary index. After integrating over the cross-section, we get:

$$\left\langle \mathbf{E}^I_p, \mathbf{H}^I_i \right\rangle + \sum_j R_{j,p} \left\langle \mathbf{E}^I_j, \mathbf{H}^I_i \right\rangle = \sum_j T_{j,p} \left\langle \mathbf{E}^{II}_j, \mathbf{H}^I_i \right\rangle \quad (2.29)$$

$$\left\langle \mathbf{E}^I_i, \mathbf{H}^I_p \right\rangle - \sum_j R_{j,p} \left\langle \mathbf{E}^I_i, \mathbf{H}^I_j \right\rangle = \sum_j T_{j,p} \left\langle \mathbf{E}^I_i, \mathbf{H}^{II}_j \right\rangle \quad (2.30)$$

where the scalar product is defined as the following overlap integral:

$$\langle \mathbf{E}_m, \mathbf{H}_n \rangle \equiv \int\int_s (\mathbf{E}_m \times \mathbf{H}_n) \cdot \mathbf{u}_z dS$$

If we decide to truncate the series expansion after N terms, we have $2N$ unknowns: N reflection coefficients and N transmission coefficients. Eq. 2.29 and 2.30 provide us exactly with $2N$ equations, since we can write them for all i in $1 \to N$. However, we can reduce the dimensionality of this linear system by invoking the orthogonality relation eq. 2.26:

$$\delta_{ip} \left\langle \mathbf{E}^I_p, \mathbf{H}^I_p \right\rangle + R_{i,p} \left\langle \mathbf{E}^I_i, \mathbf{H}^I_i \right\rangle = \sum_j T_{j,p} \left\langle \mathbf{E}^{II}_j, \mathbf{H}^I_i \right\rangle \quad (2.31)$$

$$\delta_{ip} \left\langle \mathbf{E}^I_p, \mathbf{H}^I_p \right\rangle - R_{i,p} \left\langle \mathbf{E}^I_i, \mathbf{H}^I_i \right\rangle = \sum_j T_{j,p} \left\langle \mathbf{E}^I_i, \mathbf{H}^{II}_j \right\rangle \quad (2.32)$$

Adding and subtracting these equations yields:

$$\sum_j \left[\left\langle \mathbf{E}^I_i, \mathbf{H}^{II}_j \right\rangle + \left\langle \mathbf{E}^{II}_j, \mathbf{H}^I_i \right\rangle \right] T_{j,p} = 2\delta_{ip} \left\langle \mathbf{E}^I_p, \mathbf{H}^I_p \right\rangle \quad (2.33)$$

$$R_{i,p} = \frac{1}{2 \left\langle \mathbf{E}^I_i, \mathbf{H}^I_i \right\rangle} \sum_j \left[\left\langle \mathbf{E}^{II}_j, \mathbf{H}^I_i \right\rangle - \left\langle \mathbf{E}^I_i, \mathbf{H}^{II}_j \right\rangle \right] T_{j,p} \quad (2.34)$$

This shows that we can first calculate the transmission coefficients by solving an $N \times N$ linear system, and then obtain the reflection coefficients by a simple matrix multiplication.

After obtaining R and T upon incidence of mode p, we can of course repeat the whole procedure using all modes p in $1 \to N$. Important to note is that this changes only the RHS in the linear system in eq. 2.33, so that we do not have to invert[2] another system matrix.

[2] Actually, when solving the system even an explicit inverse is not required, as we can use the LU decomposition of the system matrix.

Usually, we will choose to normalise our modes such that $\langle \mathbf{E}_i^I, \mathbf{H}_i^I \rangle = 1$. We can then write eq. 2.33 and 2.34 more compactly after defining the following overlap matrices:

$$\mathbf{O}_{I,II}(i,j) \equiv \langle \mathbf{E}_i^I, \mathbf{H}_j^{II} \rangle \tag{2.35}$$

$$\mathbf{O}_{II,I}(i,j) \equiv \langle \mathbf{E}_i^{II}, \mathbf{H}_j^I \rangle \tag{2.36}$$

With the subscript T denoting transpose, we finally get:

$$\mathbf{T}_{I,II} = 2\left(\mathbf{O}_{I,II} + \mathbf{O}_{II,I}^T\right)^{-1} \tag{2.37}$$

$$\mathbf{R}_{I,II} = \frac{1}{2}\left(\mathbf{O}_{II,I}^T - \mathbf{O}_{I,II}\right) \cdot \mathbf{T}_{I,II} \tag{2.38}$$

In these expressions $\mathbf{T}_{I,II}$ and $\mathbf{R}_{I,II}$ are the so-called transmission and reflection matrices. Their p-th columns consist of the $T_{j,p}$ and $R_{j,p}$ from eq. 2.33 and 2.34. If we collect the expansion coefficients of an arbitrary incident field in a column vector $\mathbf{A}_{inc}$, we can write very compactly for the reflected and transmitted fields:

$$\mathbf{A}_{refl} = \mathbf{R}_{I,II} \cdot \mathbf{A}_{inc} \tag{2.39}$$

$$\mathbf{A}_{trans} = \mathbf{T}_{I,II} \cdot \mathbf{A}_{inc} \tag{2.40}$$

Obviously, we can repeat the entire procedure for incidence from medium II, which gives us the matrices $\mathbf{R}_{II,I}$ and $\mathbf{T}_{II,I}$. These four matrices completely characterise the scattering that occurs at an interface.

Finally, we want to point out the similarity in structure between eq. 2.37 and 2.38 and the well-known Fresnel equations to calculate reflection and transmission for normal incidence of a plane wave upon the interface between two semi-infinite media:

$$T = \frac{2n_1}{n_1 + n_2} \tag{2.41}$$

$$R = \frac{n_1 - n_2}{n_1 + n_2} \tag{2.42}$$

In fact, eq. 2.41 and 2.42 can also be derived from the more general treatment presented here. For homogeneous layers, where the refractive index does not vary in the transverse direction, it turns out that the overlap matrices are diagonal, meaning that there is no cross-coupling between the different modes. Further evaluation of the formulas reveals that the Fresnel equations are indeed recovered.

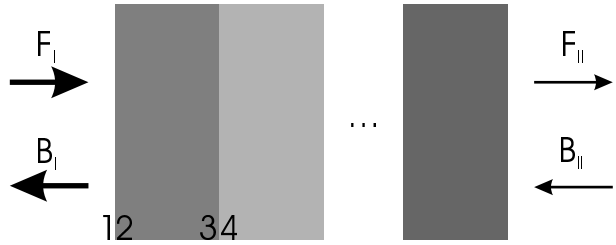

Figure 2.7: T-scheme.

2.4 Scattering at a layered stack

To calculate the reflection and transmission matrices of an entire stack, we basically need to combine two things:

- For each interface between two layers, calculate its $\mathbf{R}$ and $\mathbf{T}$-matrices, as per section 2.3.3.
- For each layer, propagate the field using the explicit knowledge of the z-dependence of the eigenmodes, given by eq. 2.9.

The precise way in which these two contributions are combined is very important, because some approaches are numerically more stable than others. In the next sections, we will discuss two of these methods: the T-scheme (based on transfer matrices) and the S-scheme (based on scattering matrices). The latter scheme is slower, but numerically more robust, and is therefore to be preferred.

2.4.1 T-scheme

Transfer matrices relate the forward and backward modes on the far side of a structure to the forward and backward modes on the near side of that structure (fig. 2.7):

$$\begin{bmatrix} \mathbf{F}_{II} \\ \mathbf{B}_{II} \end{bmatrix} = \mathbf{T} \cdot \begin{bmatrix} \mathbf{F}_{I} \\ \mathbf{B}_{I} \end{bmatrix} \tag{2.43}$$

The attractiveness of transfer matrices lies in the fact that they are easy to cascade: the $\mathbf{T}$-matrix of a structure composed from substructures is simply the product of the $\mathbf{T}$-matrices of each substructure.

To determine the transfer matrix $\mathbf{I}_{12}$ of an interface from its reflection and transmission matrices, we start from (see fig. 2.7 for the meaning of the subscripts):

$$\mathbf{F}_2 = \mathbf{T}_{12} \cdot \mathbf{F}_1 + \mathbf{R}_{21} \cdot \mathbf{B}_2 \tag{2.44}$$

$$\mathbf{B}_1 = \mathbf{R}_{12} \cdot \mathbf{F}_1 + \mathbf{T}_{21} \cdot \mathbf{B}_2 \tag{2.45}$$

These equations express that e.g. the forward field in medium 2 is the superposition of the transmitted forward field from medium 1 and the reflected backward field from medium 2. Rearranging eq. 2.44 and 2.45 into the form of eq. 2.43, yields after some matrix manipulations the matrix $\mathbf{I}_{12}$:

$$\begin{bmatrix} \mathbf{F}_2 \\ \mathbf{B}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{T}_{12} - \mathbf{R}_{21} \cdot \mathbf{T}_{21}^{-1} \cdot \mathbf{R}_{12} & \mathbf{R}_{21} \cdot \mathbf{T}_{21}^{-1} \\ -\mathbf{T}_{21}^{-1} \cdot \mathbf{R}_{12} & \mathbf{T}_{21}^{-1} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{F}_1 \\ \mathbf{B}_1 \end{bmatrix} \tag{2.46}$$

The transfer matrix $\mathbf{P}_{23}$ of a z-invariant layer is easy to write down using eq. 2.9:

$$\begin{bmatrix} \mathbf{F}_3 \\ \mathbf{B}_3 \end{bmatrix} = \begin{bmatrix} \mathbf{diag}\left(e^{-j\beta_i d_{23}}\right) & 0 \\ 0 & \mathbf{diag}\left(e^{j\beta_i d_{23}}\right) \end{bmatrix} \cdot \begin{bmatrix} \mathbf{F}_2 \\ \mathbf{B}_2 \end{bmatrix} \tag{2.47}$$

where the submatrices are diagonal matrices containing the propagation factors of each eigenmode and where d_{23} is the thickness of that layer.

After calculating the $\mathbf{I}$ and $\mathbf{P}$ matrices for the interfaces and layers, the transfer matrix for the whole stack is

$$\mathbf{T} = \mathbf{I}_{n-1,n} \cdot \ldots \cdot \mathbf{P}_{45} \cdot \mathbf{I}_{34} \cdot \mathbf{P}_{23} \cdot \mathbf{I}_{12} \tag{2.48}$$

With the knowledge of the transfer matrix of the stack, we can calculate the reflection and transmission matrices from equations 2.44 and 2.45. First we write the transfer matrix in block form:

$$\mathbf{T} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix} \tag{2.49}$$

$\mathbf{R}_{12}$ is then simply $-\mathbf{D}^{-1} \cdot \mathbf{C}$, the other matrices follow in a similar way.

The big disadvantage of the T-scheme is that for evanescent modes eq. 2.47 contains a mixture of increasing and decreasing exponentials, which is very detrimental for numerical stability. Indeed, during the calculation very small numbers will be added to very large numbers, which will obviously result in a loss of precision. In the next section, we will discuss a different scheme, where this problem does not occur.

2.4.2 S-scheme

The scattering matrix relates the outward-propagating fields of a structure to the inward-propagating fields. From eq. 2.44 and 2.45, we know that for an interface this is (see fig. 2.8 for the slightly modified subscripts)

$$\begin{bmatrix} \mathbf{F}_{1^+} \\ \mathbf{B}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{T}_{1,1^+} & \mathbf{R}_{1^+,1} \\ \mathbf{R}_{1,1^+} & \mathbf{T}_{1^+,1} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{F}_1 \\ \mathbf{B}_{1^+} \end{bmatrix} \tag{2.50}$$

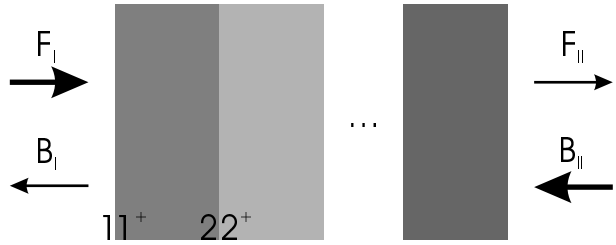

Figure 2.8: S-scheme.

In the S-scheme, we divide the stack into chunks consisting of an interface and propagation over a given distance in the exit medium of the interface [Li1996]. For e.g. the first chunk, we write the scattering matrix as

$$\begin{bmatrix} \mathbf{F}_2 \\ \mathbf{B}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{t}_{1,2} & \mathbf{r}_{2,1} \\ \mathbf{r}_{1,2} & \mathbf{t}_{2,1} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{F}_1 \\ \mathbf{B}_2 \end{bmatrix} \tag{2.51}$$

The chunk matrices $\mathbf{r}$ and $\mathbf{t}$ follow easily by combining eq. 2.50 and 2.47:

$$\mathbf{t}_{1,2} = \mathbf{diag}\left(e^{-j\beta_i d_{12}}\right) \cdot \mathbf{T}_{1,1^+} \tag{2.52}$$

$$\mathbf{r}_{2,1} = \mathbf{diag}\left(e^{-j\beta_i d_{12}}\right) \cdot \mathbf{R}_{1^+,1} \cdot \mathbf{diag}\left(e^{-j\beta_i d_{12}}\right) \tag{2.53}$$

$$\mathbf{r}_{1,2} = \mathbf{R}_{1,1^+} \tag{2.54}$$

$$\mathbf{t}_{2,1} = \mathbf{T}_{1^+,1} \cdot \mathbf{diag}\left(e^{-j\beta_i d_{12}}\right) \tag{2.55}$$

We can write such expressions for any chunk in the stack. Note that only exponentials with the same sign occur in these equations, which makes them numerically more stable than the T-scheme.

Now we only need to determine how to calculate the scattering matrix of the entire stack, given the scattering matrices of the individual chunks. More specifically, we will do this by calculating S for the first p chunks. given that S for the first $p-1$ chunks is already known. So, the following quantities are given:

$$\begin{bmatrix} \mathbf{F}_p \\ \mathbf{B}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{T}_{1,p} & \mathbf{R}_{p,1} \\ \mathbf{R}_{1,p} & \mathbf{T}_{p,1} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{F}_1 \\ \mathbf{B}_p \end{bmatrix} \tag{2.56}$$

We want to calculate the matrices

$$\begin{bmatrix} \mathbf{F}_{p+1} \\ \mathbf{B}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{T}_{1,p+1} & \mathbf{R}_{p+1,1} \\ \mathbf{R}_{1,p+1} & \mathbf{T}_{p+1,1} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{F}_1 \\ \mathbf{B}_{p+1} \end{bmatrix} \tag{2.57}$$

We do this by combining eq. 2.56 with the expressions for the chunk matrices in chunk p. After some straightforward but rather tedious algebra, we eventually arrive at ($\mathbf{I}$ is the unit matrix)

$$\mathbf{T}_{1,p+1} = \mathbf{t}_{p,p+1} \cdot \left(\mathbf{I} - \mathbf{R}_{p,1} \cdot \mathbf{r}_{p,p+1}\right)^{-1} \cdot \mathbf{T}_{1,p} \quad (2.58)$$

$$\mathbf{R}_{p+1,1} = \mathbf{t}_{p,p+1} \cdot \left(\mathbf{I} - \mathbf{R}_{p,1} \cdot \mathbf{r}_{p,p+1}\right)^{-1} \cdot \mathbf{R}_{p,1} \cdot \mathbf{t}_{p+1,p} + \mathbf{r}_{p+1,p} \quad (2.59)$$

$$\mathbf{R}_{1,p+1} = \mathbf{T}_{p,1} \cdot \left(\mathbf{I} - \mathbf{r}_{p,p+1} \cdot \mathbf{R}_{p,1}\right)^{-1} \cdot \mathbf{r}_{p,p+1} \cdot \mathbf{T}_{1,p} + \mathbf{R}_{1,p} \quad (2.60)$$

$$\mathbf{T}_{p+1,1} = \mathbf{T}_{p,1} \cdot \left(\mathbf{I} - \mathbf{r}_{p,p+1} \cdot \mathbf{R}_{p,1}\right)^{-1} \cdot \mathbf{t}_{p+1,p} \quad (2.61)$$

Using these formulas, we can calculate the scattering matrix of the entire stack, starting from the first chunk (eq. 2.51) and working our way through the rest of the chunks.

We already mentioned that only decreasing exponentials occur in these formulas. More rigorous studies [Li1996] show that the numerical stability of the S-scheme is indeed far better than that of the T-scheme. However, this stability comes at a price, because more matrix multiplications have to be performed in the S-scheme.

Finally, we would also like to mention the existence of a third scheme, the R-scheme, which relates the electric and magnetic fields on both sides of a structure. Such a scheme is popular in the Method of Lines [Pregla1996], a semi-numerical eigenmode method which employs spatial discretisation in the transverse direction. However, the R-scheme seems to have an even higher numerical overhead [Li1996].

2.4.3 Finite periodic stacks

In many cases we have to calculate the scattering matrix of a stack that consists of a finite number of repetitions of the same basic period. This is e.g. the case when dealing with Bragg mirrors or other photonic crystal structures. In eq. 2.58 through 2.61, we have seen how to concatenate different substacks using the S-scheme formalism. If we define the concatenation of two substacks A and B by the operation $A \otimes B$, we could calculate a stack S_N consisting of N periodic repetitions of S by performing the following concatenation:

$$S_N = S \otimes S \otimes ... \otimes \mathrm{S} \ (N \text{ times}) \quad (2.62)$$

The computational cost associated with this expression is linear in the number of periods. However, we can do better than that, if we make use of the following recurrence relations:

$$S_{2i} = S_i \otimes S_i \quad (2.63)$$

$$S_{2i+1} = S_{2i} \otimes S \quad (2.64)$$

This is the same principle as calculating e.g. a^4 as $\left(a^2\right)^2$, rather than as $aaaa$. The former expression requires only two multiplications, whereas the latter requires four.

For practical purposes, such a recursive formulation requires much more memory, since all the intermediate results have to be stored. This can easily be avoided by formulating this scheme iteratively, rather than recursively. We can do this by using the binary representation N_b of the number of periods, as illustrated in the following pseudocode:

```
result ← S
loop from second bit in N_b to last bit:
    result ← result ⊗ result
    if current bit in N_b equals 1
        result ← result ⊗ S
```

In this way, the computation time is only logarithmic in the number of periods, rather than linear. This can mean substantial savings, especially for large numbers of periods.

Also note that we do not sacrifice numerical stability in any way. In chapter 5, we will touch upon another technique to perform periodic extension, based on the eigenvectors of the transfer matrix, but this approach is numerically less robust and is therefore to be avoided.

2.5 Fields inside a stack

The scattering matrix of a stack gives us information about the field (or rather its expansion in eigenmodes) on the near and far edges of the stack. Sometimes it is also useful to know the fields at a position inside the stack, given a certain incident field. We will only consider ourselves here with interior positions at the interface between two layers, since the fields inside a layer can easily be propagated from the fields at the surrounding interfaces using eq. 2.9.

Suppose that we excite a stack with an incident field $\mathbf{F}_1$ from the left, and that no exterior field $\mathbf{B}_3$ is coming from the right (fig. 2.9). The presence of $\mathbf{F}_1$ will give rise to a reflected field $\mathbf{B}_1 = \mathbf{R}_{13} \cdot \mathbf{F}_1$, with $\mathbf{R}_{13}$ the reflection matrix of the entire stack. To calculate the internal fields $\mathbf{F}_2$ and $\mathbf{B}_2$, we divide the stack S_{13} in two substacks S_{12} and S_{23}, for which we can calculate the reflection and transmission matrices. We can now proceed in two ways, depending on whether we use a T or an S-like formulation for the incident fields. Note that this has no bearing on the scheme used to calculate the scattering matrices, which is preferably the S-scheme. We will discuss both formulations for the incident fields and discuss their numerical advantages and disadvantages.

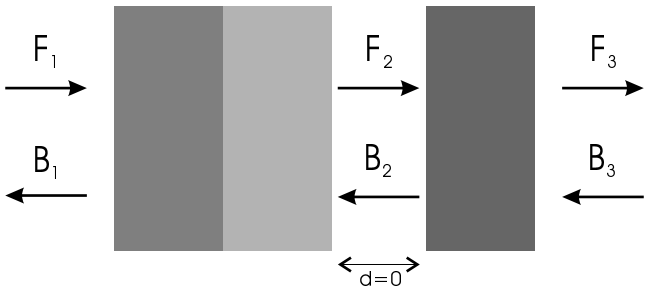

Figure 2.9: Fields inside a stack.

2.5.1 T-excitation

Eq. 2.46 describing the transfer matrix immediately yields for the unknown fields that

$$\begin{bmatrix} \mathbf{F}_2 \\ \mathbf{B}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{T}_{12} - \mathbf{R}_{21} \cdot \mathbf{T}_{21}^{-1} \cdot \mathbf{R}_{12} & \mathbf{R}_{21} \cdot \mathbf{T}_{21}^{-1} \\ -\mathbf{T}_{21}^{-1} \cdot \mathbf{R}_{12} & \mathbf{T}_{21}^{-1} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{F}_1 \\ \mathbf{R}_{13} \cdot \mathbf{F}_1 \end{bmatrix} \tag{2.65}$$

For a position deep inside the stack however, $\mathbf{T}_{21}$ can become numerically close to singular. This is due to higher order modes, which will be strongly damped when propagating and which will therefore have near zero transmission through S_{12}. This introduces numbers close to zero on the diagonal of the otherwise diagonally dominant matrix $\mathbf{T}_{21}$. In chapter 5 we will discuss a "best effort"-attempt at inverting this matrix, but because this method is only an approximation, it is not recommended when full precision is required.

2.5.2 S-excitation

An alternative is to start from the scattering matrix equation 2.44 and to also involve the right-hand part of the stack in our discussion:

$$\mathbf{F}_2 = \mathbf{T}_{12} \cdot \mathbf{F}_1 + \mathbf{R}_{21} \cdot \mathbf{B}_2 \tag{2.66}$$

$$\mathbf{B}_2 = \mathbf{R}_{23} \cdot \mathbf{F}_2 + \mathbf{T}_{32} \cdot \mathbf{B}_3 \tag{2.67}$$

In the absence of an incident field $\mathbf{B}_3$ in medium 3, we can therefore write

$$\mathbf{F}_2 = \left[\mathbf{I} - \mathbf{R}_{21} \cdot \mathbf{R}_{23}\right]^{-1} \cdot \mathbf{T}_{21} \cdot \mathbf{F}_1 \tag{2.68}$$

$$\mathbf{B}_2 = \mathbf{R}_{23} \cdot \mathbf{F}_2 \tag{2.69}$$

In this way we avoid inverting $\mathbf{T}_{21}$. The numerical effort is however higher, because we also have to calculate the scattering matrix of the right-hand part of the stack.

It is instructive to take a closer look at the matrix $\mathbf{I} - \mathbf{R}_{21} \cdot \mathbf{R}_{23}$, that we have to invert in eq. 2.68, because similar terms also figure in the expressions for the S-scheme (eq. 2.58 through 2.61). This matrix will become singular if $\mathbf{R}_{21} \cdot \mathbf{R}_{23}$ equals the unit matrix, which means that a field first reflecting off the right stack and then off the left stack will reproduce itself. This condition of unit round trip gain is precisely the one that determines lasing (or eigenmode) resonances of the entire stack. When this happens, the expressions 2.68 and 2.69 will become singular and cannot be used. However, this is rarely a problem, because in these cases we typically first determine the fields in the middle of the stack[3], and then propagate them through either the left or right substack, which by themselves will not exhibit these resonances.

As a final note in this chapter, we want to point out that because of all the matrix manipulations involved, the calculation time of eigenmode methods roughly scales as N^3, with N the number of modes retained in the series expansion. Any technique therefore that allows us to obtain results with the same accuracy, but with a smaller number of modes, will present significant time savings. In the next chapter, some methods will be presented that do precisely that.

Appendix A: Slab waveguides

We will now fill in the final missing pieces of the puzzle by providing the geometry-dependent information that is needed for the eigenmode expansion: the dispersion relations that determine the propagation constants, the field profiles of the modes and the overlap integrals between two modes. In this appendix, we will do this for arbitrary slab waveguides, while in the next appendix, we will deal with circular structures consisting of a central core and an outer cladding. For the boundary conditions, we will choose the simple perfectly conducting electric wall, and leave more sophisticated boundaries for chapter 3. In order to avoid an overly large number of equations, we will focus on the broad outline of the derivation and leave the intermediate steps for the interested reader to verify.

A.1 Dispersion relation

Consider a slab waveguide consisting of an arbitrary number of layers with refractive index n_i and thickness d_i. The structure is two-dimensional, i.e. the geometry does not vary in the y-direction. These

[3] See also chapter 6.

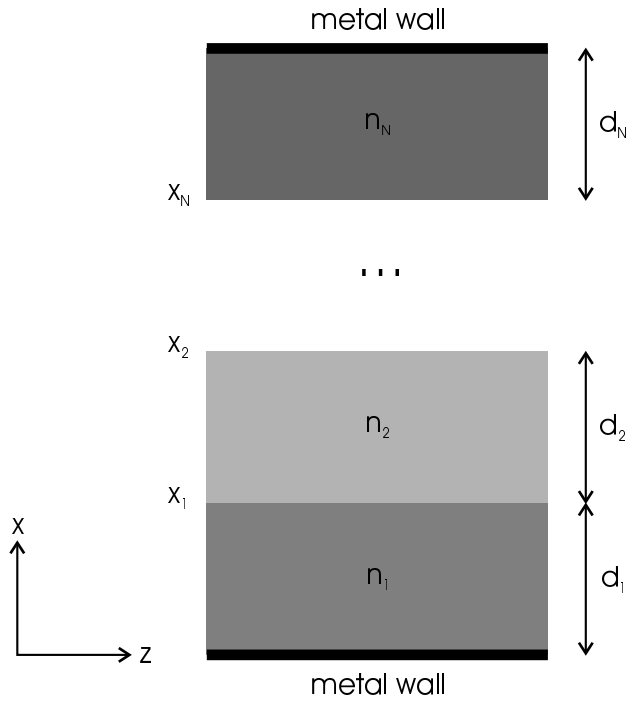

Figure 2.10: Slab waveguide.

layers are placed between two metal walls that force the tangential component of the electric field to zero (fig. 2.10). We will derive the dispersion relation for TE modes, the one for TM modes follows similarly. If we assume a given (but so far unknown) propagation constant β for an eigenmode, the transverse components of the wavevector in each layer i are then determined by [Schlereth1990]

$$k_{t,i} = \sqrt{\left(\frac{2\pi}{\lambda_0} n_i\right)^2 - \beta^2} \tag{2.70}$$

Here, λ_0 is the vacuum wavelength of the light. The electric field in each layer can be written as a sum of forward and backward propagating plane waves, with a wavevector as per eq. 2.70:

$$\mathbf{E}_i(x,y,z) = \mathbf{u}_y E_i(x)\, e^{-j\beta z} \tag{2.71}$$

$$E_i = F_i e^{-jk_{t,i}(x-x_i)} + B_i e^{jk_{t,i}(x-x_i)} \tag{2.72}$$

We can now relate the fields in subsequent waveguides with the T-formalism. For the scattering matrix of the interface, we use the

Fresnel formulas [Lee1986], while the propagation matrix is once again given by eq. 2.47. Combined, these give eventually [Smith1991]

$$\begin{bmatrix} F_i \\ B_i \end{bmatrix} = \frac{1}{2} \begin{bmatrix} (1 + k_{t,i}/k_{t,i-1})\, e^{-jk_{t,i}d_i} & (1 - k_{t,i}/k_{t,i-1})\, e^{jk_{t,i}d_i} \\ (1 - k_{t,i}/k_{t,i-1})\, e^{-jk_{t,i}d_i} & (1 + k_{t,i}/k_{t,i-1})\, e^{jk_{t,i}d_i} \end{bmatrix} \cdot \begin{bmatrix} F_{i-1} \\ B_{i-1} \end{bmatrix} \tag{2.73}$$

To determine the equation that will give us the values of β that satisfy the boundary conditions, we start out by choosing $F_0 = 1$ and $B_0 = -1$ at one boundary, such that $F_0 + B_0 = 0$. Subsequently, we calculate F_N and B_N at the other boundary using eq. 2.73 and impose that there too $F_N + B_N = 0$.

Because we are only interested in the ratio of F and B, we can factor out and discard the increasing exponential $e^{jk_{t,i}d_i}$ from the matrix in eq. 2.73, such that only decreasing exponentials remain. In this way, we avoid numerical instabilities associated with mixing large and small numbers.

A.2 Field profiles

Calculating field profiles can be done from eq. 2.73: start by choosing $F_0 = 1$ and $B_0 = -1$ at the lower boundary and subsequently work through all the layers. However, for numerical stability is a preferable to use the S-scheme rather than the T-scheme for calculating field profiles.

A.3 Overlap integrals

Because of the simple structure of eq. 2.72, the overlap integrals can be calculated analytically very easily. They all involve integrals similar to

$$\int_{x_1}^{x_2} e^{-jk_{x,I}x} . e^{-jk_{x,II}x} dx \tag{2.74}$$

which of course evaluates as

$$\frac{j}{k_{x,I} + k_{x,II}} \left[e^{-jk_{x,I}x} . e^{-jk_{x,II}x} \right]_{x_1}^{x_2} \quad \text{if} \quad k_{x,I} + k_{x,II} \neq 0 \tag{2.75}$$

$$x_2 - x_1 \quad \text{if} \quad k_{x,I} + k_{x,II} = 0 \tag{2.76}$$

Appendix B: Circular waveguides

Fig. 2.11 shows a circular waveguide, consisting of a core with refractive index n_{co} and radius r, and a cladding with index n_{cl}. The whole structure is enclosed in a perfectly conducting cylinder with radius R.

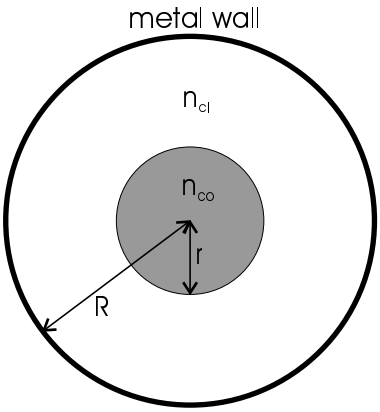

Figure 2.11: Circular waveguide.

B.1 Field profiles

The nabla operator in cylindrical coordinates takes the following form:

$$\nabla = \frac{\partial}{\partial \rho}\mathbf{u}_\rho + \frac{1}{\rho}\frac{\partial}{\partial \phi}\mathbf{u}_\phi + \frac{\partial}{\partial z}\mathbf{u}_z \tag{2.77}$$

Inserting this into Maxwell's equations for z-invariant media ultimately leads to the following equation for E_z and H_z [Vassallo1985]:

$$\left\{ \frac{\partial^2}{\partial \rho^2} + \frac{1}{\rho}\frac{\partial}{\partial \rho} + \frac{1}{\rho^2}\frac{\partial^2}{\partial \phi^2} + k^2 - \beta^2 \right\} \left\{ \begin{array}{c} E_z \\ H_z \end{array} \right\} = 0 \tag{2.78}$$

The general solution to this equation in any uniform medium is

$$\left(AJ_n\left(k_t\rho\right) + BH_n^{(2)}\left(k_t\rho\right) \right) \left(Ccos\left(n\phi\right) + Dsin\left(n\phi\right) \right) \tag{2.79}$$

Here, $k_t = \sqrt{k^2 - \beta^2}$ and n is an integer describing the order of the Bessel function J and of the Hankel function of the second kind $H^{(2)}$. In fact, J can be seen as the cylindrical equivalent of the $\sin$ or $\cos$ functions, while $H^{(2)}$ acts like an outgoing wave e^{-jkr}. The solutions for different n are mutually orthogonal, as are those for $\cos\left(n\phi\right)$ and $\sin\left(n\phi\right)$.

In what follows, we will take the solution of E_z proportional to $\cos\left(n\phi\right)$, and that of H_z proportional to $\sin\left(n\phi\right)$. The dual case follows similarly. Making abstraction of the angular dependence for the moment, we can propose the following behaviour for E_z and H_z in the core and the cladding:

$$E_z\left(\rho\right) = A\frac{J_n\left(k_{t,co}\rho\right)}{J_n\left(k_{t,co}r\right)},\ \rho < r \tag{2.80}$$

$$= A\frac{J_n\left(k_{t,cl}R\right)H_n^{(2)}\left(k_{t,cl}\rho\right) - H_n^{(2)}\left(k_{t,cl}R\right)J_n\left(k_{t,cl}\rho\right)}{J_n\left(k_{t,cl}R\right)H_n^{(2)}\left(k_{t,cl}r\right) - H_n^{(2)}\left(k_{t,cl}R\right)J_n\left(k_{t,cl}r\right)},\ \rho > r$$

$$H_z\left(\rho\right) = B\frac{J_n\left(k_{t,co}\rho\right)}{J_n\left(k_{t,co}r\right)},\ \rho < r \tag{2.81}$$

$$= B\frac{J_n^{'}\left(k_{t,cl}R\right)H_n^{(2)}\left(k_{t,cl}\rho\right) - H_n^{(2)'}\left(k_{t,cl}R\right)J_n\left(k_{t,cl}\rho\right)}{J_n^{'}\left(k_{t,cl}R\right)H_n^{(2)}\left(k_{t,cl}r\right) - H_n^{(2)'}\left(k_{t,cl}R\right)J_n\left(k_{t,cl}r\right)},\ \rho > r$$

These expressions automatically take care of the continuity of E_z and H_z at $\rho = r$. Also, the tangential electric field is identically zero at the wall. In the core, no Hankel function can be present, because it is singular at the origin. Note that the primes in eq. 2.81 denote derivation with respect to the whole argument. A and B are currently unknown, but their ratio will follow from the dispersion relation, as will of course the propagation constant β.

With the knowledge of the z-components of the fields, we can also derive the other field components according to [Vassallo1985]

$$E_\rho = \frac{j}{k_t^2}\left(-\frac{\omega\mu}{\rho}\frac{\partial H_z}{\partial\phi} - \beta\frac{\partial E_z}{\partial\rho}\right) \tag{2.82}$$

$$E_\phi = \frac{j}{k_t^2}\left(\omega\mu\frac{\partial H_z}{\partial\rho} - \frac{\beta}{\rho}\frac{\partial E_z}{\partial\phi}\right) \tag{2.83}$$

$$H_\rho = \frac{j}{k_t^2}\left(\frac{\omega\varepsilon}{\rho}\frac{\partial E_z}{\partial\phi} - \beta\frac{\partial H_z}{\partial\rho}\right) \tag{2.84}$$

$$H_\phi = \frac{j}{k_t^2}\left(-\omega\varepsilon\frac{\partial E_z}{\partial\rho} - \frac{\beta}{\rho}\frac{\partial H_z}{\partial\phi}\right) \tag{2.85}$$

B.2 Dispersion relation

Imposing the continuity of E_ϕ and H_ϕ at $\rho = r$ leads to

$$B\omega\mu_0\left(\frac{\mu_{r,cl}}{k_{t,cl}}\frac{a_{TE}}{b_{TE}} - \frac{\mu_{r,co}}{k_{t,co}}\frac{J_n^{'}\left(k_{t,co}r\right)}{J_n\left(k_{t,co}r\right)}\right) = A\frac{\beta n}{r}\left(\frac{1}{k_{t,co}^2} - \frac{1}{k_{t,cl}^2}\right) \tag{2.86}$$

$$A\omega\varepsilon_0\left(\frac{\varepsilon_{r,cl}}{k_{t,cl}}\frac{a_{TM}}{b_{TM}} - \frac{\varepsilon_{r,co}}{k_{t,co}}\frac{J_n^{'}\left(k_{t,co}r\right)}{J_n\left(k_{t,co}r\right)}\right) = B\frac{\beta n}{r}\left(\frac{1}{k_{t,co}^2} - \frac{1}{k_{t,cl}^2}\right) \tag{2.87}$$

with

$$a_{TE} \equiv J_n^{'}\left(k_{t,cl}R\right)H_n^{(2)'}\left(k_{t,cl}r\right) - H_n^{(2)'}\left(k_{t,cl}R\right)J_n^{'}\left(k_{t,cl}r\right) \tag{2.88}$$

$$b_{TE} \equiv J_n^{'}\left(k_{t,cl}R\right)H_n^{(2)}\left(k_{t,cl}r\right) - H_n^{(2)'}\left(k_{t,cl}R\right)J_n\left(k_{t,cl}r\right) \tag{2.89}$$

$$a_{TM} \equiv J_n\left(k_{t,cl}R\right) H_n^{(2)'}\left(k_{t,cl}r\right) - H_n^{(2)}\left(k_{t,cl}R\right) J_n'\left(k_{t,cl}r\right) \quad (2.90)$$

$$b_{TM} \equiv J_n\left(k_{t,cl}R\right) H_n^{(2)}\left(k_{t,cl}r\right) - H_n^{(2)}\left(k_{t,cl}R\right) J_n\left(k_{t,cl}r\right) \quad (2.91)$$

Eliminating A and B from eq. 2.86 and 2.87 gives us the required dispersion relation:

$$\left(\frac{\mu_{r,cl}}{k_{t,cl}}\frac{a_{TE}}{b_{TE}} - \frac{\mu_{r,co}}{k_{t,co}}\frac{J_n'\left(k_{t,co}r\right)}{J_n\left(k_{t,co}r\right)}\right)\left(\frac{\varepsilon_{r,cl}}{k_{t,cl}}\frac{a_{TM}}{b_{TM}} - \frac{\varepsilon_{r,co}}{k_{t,co}}\frac{J_n'\left(k_{t,co}r\right)}{J_n\left(k_{t,co}r\right)}\right) = \frac{\beta^2 n^2}{r^2}\left(\frac{1}{k_{t,co}^2} - \frac{1}{k_{t,cl}^2}\right)^2 \quad (2.92)$$

Contrary to slab waveguides, the modes generally cannot be separated into TE and TM, due to the cross-coupling that occurs at the interface. The only exception is the case $n = 0$, where the dispersion relation factors into a TE and TM part.

B.3 Overlap integrals

The overlap integral to be calculated has the following form:

$$\int\int_S \left(\mathbf{E}_I \times \mathbf{H}_{II}\right) \cdot \mathbf{u}_z dS = \int\int_S \left(E_{\rho,I}H_{\phi,II} - H_{\rho,II}E_{\phi,I}\right) \rho d\rho d\phi \quad (2.93)$$

The angular argument runs from 0 to 2π, while the radial domain can be split up in parts where the material properties do not change for both media. From eq. 2.82 through 2.85, we can gather that in these expression, two types of integrals will figure:

$$I_1 \equiv \int\int \left(\frac{1}{\rho}\frac{\partial A}{\partial \rho}\frac{\partial B}{\partial \phi} - \frac{1}{\rho}\frac{\partial A}{\partial \phi}\frac{\partial B}{\partial \rho}\right) \rho d\rho d\phi \quad (2.94)$$

$$I_2 \equiv \int\int \left(\frac{\partial A}{\partial \rho}\frac{\partial B}{\partial \rho} - \frac{1}{\rho^2}\frac{\partial A}{\partial \phi}\frac{\partial B}{\partial \phi}\right) \rho d\rho d\phi \quad (2.95)$$

where A and B stand in for either E_z or H_z and are linear combinations of Bessel and Hankel functions. If we assume for A a cosine angular dependence and for B a sine dependence, the first integral becomes

$$I_1 = n\int_0^{2\pi} d\phi \int_{\rho_1}^{\rho 2} \left(\cos^2\left(n\phi\right)\frac{\partial A}{\partial \rho}B + \sin^2\left(n\phi\right)A\frac{\partial B}{\partial \rho}\right) d\rho \quad (2.96)$$

$$= n\pi \int_{\rho_1}^{\rho_2} \frac{\partial}{\partial \rho}\left(AB\right) d\rho \quad (2.97)$$

$$= n\pi \left[AB\right]_{\rho_1}^{\rho_2} \quad (2.98)$$

The second integral is less obvious to calculate. It is useful to derive an auxiliary relation for Bessel functions first. We start from the Bessel differential equations eq. 2.78:

$$\frac{1}{\rho}\frac{\partial}{\partial\rho}\left(\rho\frac{\partial A}{\partial\rho}\right)-\left(\frac{n^2}{\rho^2}-k_{t,A}^2\right)A = 0 \tag{2.99}$$

$$\frac{1}{\rho}\frac{\partial}{\partial\rho}\left(\rho\frac{\partial B}{\partial\rho}\right)-\left(\frac{n^2}{\rho^2}-k_{t,B}^2\right)B = 0 \tag{2.100}$$

Multiplying eq. 2.99 by B, eq. 2.100 by A, substracting both equations and integrating by parts gives

$$\int AB\rho d\rho = \frac{\rho}{k_{t,A}^2-k_{t,B}^2}\left[A\frac{\partial B}{\partial\rho}-B\frac{\partial A}{\partial\rho}\right] \tag{2.101}$$

Coming back to our integral I_2, we can write it after introducing the angular dependencies like this:

$$I_2 = \pi\int\frac{\partial A}{\partial\rho}\frac{\partial B}{\partial\rho}\rho d\rho+\pi n^2\int\frac{AB}{\rho^2}\rho d\rho \tag{2.102}$$

Integrating the first term by parts yields

$$I_2 = \pi A\frac{\partial B}{\partial\rho}\rho-\pi\int A\left[\frac{\partial^2 B}{\partial\rho^2}+\frac{1}{\rho}\frac{\partial B}{\partial\rho}\right]\rho d\rho+\pi n^2\int\frac{AB}{\rho^2}\rho d\rho \tag{2.103}$$

Combining the second and the third term and using the Bessel differential equation 2.100 now gives

$$I_2 = \pi A\frac{\partial B}{\partial\rho}\rho+\pi k_{t,B}^2\int AB\rho d\rho \tag{2.104}$$

This last integral can be solved by using eq. 2.101 to give

$$I_2 = \frac{\pi\rho}{k_{t,A}^2-k_{t,B}^2}\left(k_{t,A}^2A\frac{\partial B}{\partial\rho}-k_{t,B}^2B\frac{\partial A}{\partial\rho}\right) \tag{2.105}$$

In the case where $k_{t,A}=k_{t,B}$, like e.g. when normalising a mode, we can follow a similar line of reasoning, but this time by using a modified version of eq. 2.101, obtained after taking the limit $k_{t,B}\to k_{t,A}$ with de L'Hôpital's rule:

$$\int A^2\rho d\rho = \frac{\rho^2}{2}\left[\left(1-\frac{n^2}{\rho^2k_{t,A}^2}\right)A^2+\frac{1}{k_{t,A}^2}\left(\frac{\partial A}{\partial\rho}\right)^2\right] \tag{2.106}$$

Chapter 3

Boundary conditions

"It's got to be perfect"
Fairground Attraction

3.1 Introduction

Choosing the right boundary conditions is crucially important for the accuracy and efficiency of the eigenmode expansion method, or indeed any numerical method, which is what warrants their treatment in a separate chapter. We will illustrate the principles of several different boundary conditions using the TE modes of the multi-slab waveguide. The treatment of TM modes follows along similar lines, as does that of circular waveguides. We will try to present a unified treatment of the different boundary conditions, and therefore we will first derive the dispersion relation for a slab waveguide between two generalised boundaries.

Consider the situation depicted in fig. 3.1. The actual waveguide structure will be treated as a black box, described by its transfer matrix. The boundary at $x = 0$ will impose a certain relation between forward and backward-propagating waves, e.g. described by a reflection coefficient $r_0 = F_0/B_0$. Here F relates to the forward-propagating field in the $+x$ direction (i.e. away from the boundary), and B to the backward field propagating along $-x$. We can therefore write this boundary condition as

$$F_0 - r_0 B_0 = 0 \tag{3.1}$$

For the TE case, F and B are directly proportional to the electric field E_y (see appendix A of the previous chapter). A field distribution that satisfies this condition is trivially $F_0 = r_0$ and $B_0 = 1$. Using the transfer matrix formalism, we can propagate this field through the lower

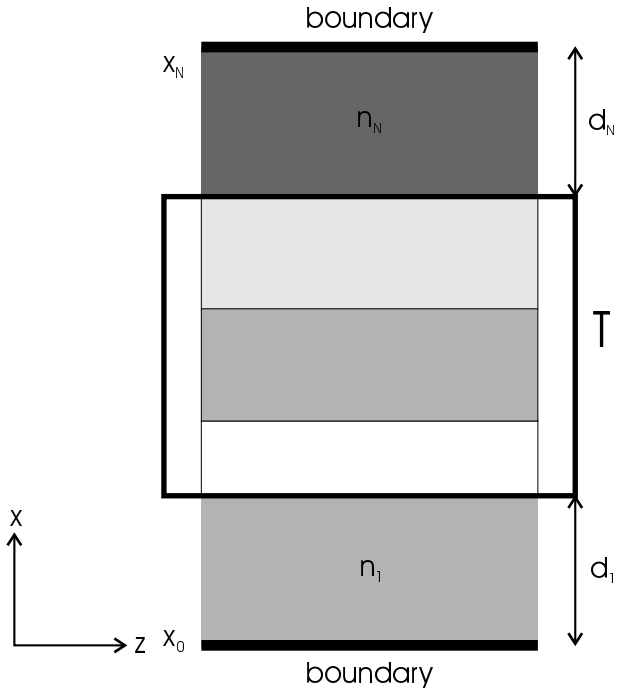

Figure 3.1: Slab waveguide with general boundaries.

cladding, the black box and the upper cladding to arrive at the second boundary. At this wall, we impose a similar boundary condition:

$$r_N F_N - B_N = 0 \tag{3.2}$$

Note that in this case the reflected field propagates in the $-x$ direction. Expressing that the field at the upper boundary satisfies eq. 3.2, gives us the dispersion relation for a slab waveguide with generalised boundary conditions:

$$\begin{bmatrix} r_N e^{-jd_N k_{x,N}} & -e^{+jd_N k_{x,N}} \end{bmatrix} \cdot \begin{bmatrix} T_{11} & T_{12} \\ T_{21} & T_{22} \end{bmatrix} \cdot \begin{bmatrix} r_0 e^{-jd_1 k_{x,1}} \\ e^{+jd_1 k_{x,1}} \end{bmatrix} = 0 \tag{3.3}$$

In the following sections, we will discuss different boundary conditions that can be expressed in this way. Note that in the most general case, the reflection coefficients r_0 and r_N can be functions of k_x. Stated otherwise, this means that the boundary conditions can depend on the incidence angle. We will concentrate on the boundary at $x = 0$, the other one obviously being completely analogous.

3.2 Hard walls

3.2.1 PEC and PMC walls

The most straightforward boundary conditions to use are perfect electric or magnetic conductors (PECs or PMCs). An electric wall imposes that the total tangential **E**-field vanishes at the boundary, which corresponds to the well-known Dirichlet boundary condition:

$$E_y(x = 0) = Fe^{-jk_x x} + Be^{+jk_x x} = 0 \tag{3.4}$$

Demanding that $F + B = 0$ amounts to choosing $r = -1$.

A magnetic wall imposes a zero tangential **H**-field. We can easily derive from Maxwell's equations that this means for TE modes that

$$H_z(x = 0) = \frac{j}{k_0 c\mu}\frac{dE_y}{dx} = \frac{k_x}{k_0 c\mu}\left(Fe^{-jk_x x} - Be^{+jk_x x}\right) = 0 \tag{3.5}$$

So, for a PMC we choose $r = 1$. Imposing a vanishing derivative of E_y also means that this is a von Neumann boundary condition.

Both electric and magnetic walls are what is commonly referred to as 'hard' walls: they reflect all the incident power, because $|r| = 1$. This can be problematic, especially when we want to model structures with high radiation losses, i.e. structures that emit radiation that would otherwise propagate freely towards infinity. The presence of a PEC or a PMC wall will send this radiation back to the structure that we want to study, and this reflected power can seriously compromise the

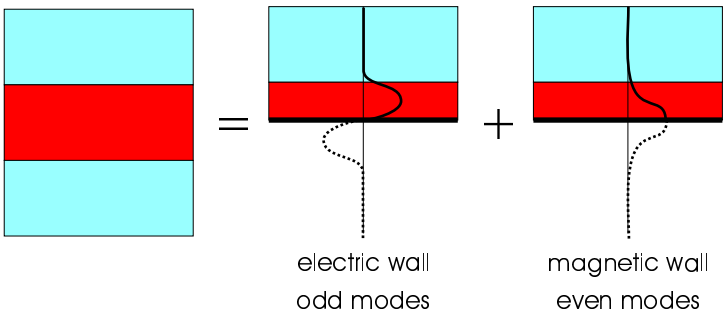

Figure 3.2: The use of PEC and PMC walls to exploit symmetry.

accuracy of the obtained results. We will show an example illustrating this problem in the next section.

A use of PEC and PMC walls that is not plagued by parasitic reflections is in exploiting the symmetry of a structure in order to reduce the computational effort. This is illustrated in fig. 3.2. The waveguide there has a mirror symmetry along the central axis. Such a structure contains even and odd modes, which do not couple among themselves (i.e. the overlap integral of an even mode with an odd mode is always zero). By placing either an electric or a magnetic wall at the symmetry axis, we can effectively cut the structure in two. Therefore, we can separate the problem into two non-interacting subproblems, involving only even or only odd modes. In this way, we reduce the computation time from $(2N)^3$ to $N^3 + N^3$.

3.2.2 Parasitic reflections

To illustrate the nefarious effects of parasitic reflections from hard walls, we model the 1D grating (or 1D photonic crystal) from fig. 3.3. This structure was the subject of a modelling exercise within the European COST268 action. It consists of a Si_3N_4 ($n = 2$) waveguide on a SiO_2 ($n = 1.45$) substrate, with a 20-period grating etched into it. The top cladding is air, with a metal wall placed at a distance d from the waveguide. The waveguide is 500 nm thick, while the etch depth is 125 nm. The period length is 430 nm, and the filling factor of the grating is 50%. We excite the structure from the left with the fundamental mode of the unetched waveguide. Part of this incident light is reflected back into the waveguide (R), part of it is transmitted to the fundamental mode of the exit waveguide (T), while the rest of the power is lost to radiation and scattering (L). The thickness of the substrate layer between the waveguide and the bottom wall is chosen to be $1.45d$, to have a similar optical thickness as the air cladding. The number of modes used in these simulations is $N = 120/\lambda * d/5$. The convergence

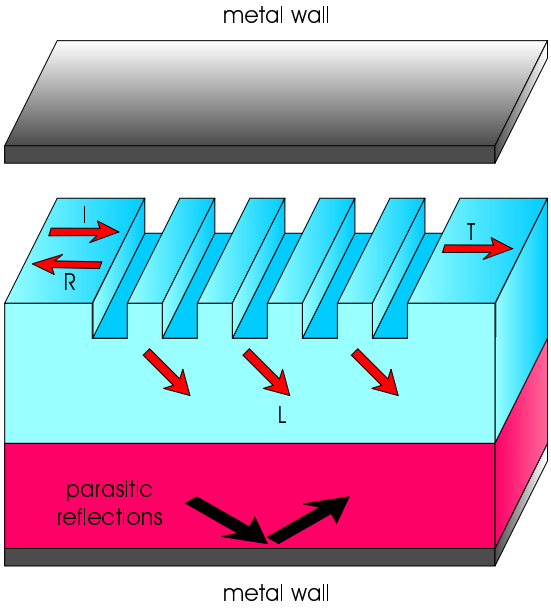

Figure 3.3: Waveguide with a shallow grating etched into it.

as a function of N will be illustrated in section 3.5, but the number of modes in this example is large enough to be independent of a further increase in N.

Fig. 3.4 shows the reflected power R as a function of the wavelength, for different thicknesses d of the air layer between the waveguide and the top electric wall. We see that the grating acts like a Bragg mirror around the design wavelength of 1.55 μm, with a reflectivity of around 50%. The different curves in the figure overlap very well around this wavelength, meaning that there is no influence from parasitic reflections. For short wavelengths however, the curves start to exhibit spikes, and the spike maxima are at different wavelengths for different values of d. This behaviour is even more prominent if we plot the loss L of the mode in fig. 3.5, defined as $L = 1 - R - T$.

It is very clear that the presence of the perfectly reflecting boundary conditions disturbs the results. To illustrate the seriousness of this problem ever further, fig. 3.6 shows the loss at a wavelength of 850 nm, as a function of the distance d to the metal wall. This plot is strongly oscillating, and it is very difficult to predict to what final value the loss will converge.

This is a serious issue, that could potentially threaten the applicability of the eigenmode expansion method to this kind of problems. It is therefore important to introduce better boundary conditions that do not reflect all the incident power.

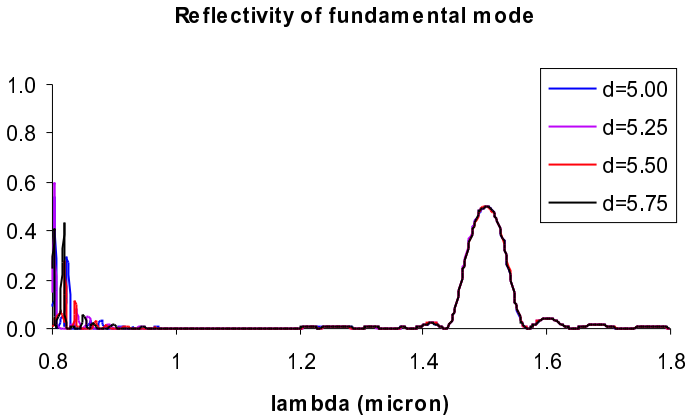

Figure 3.4: Reflectivity of the fundamental mode of the waveguide grating

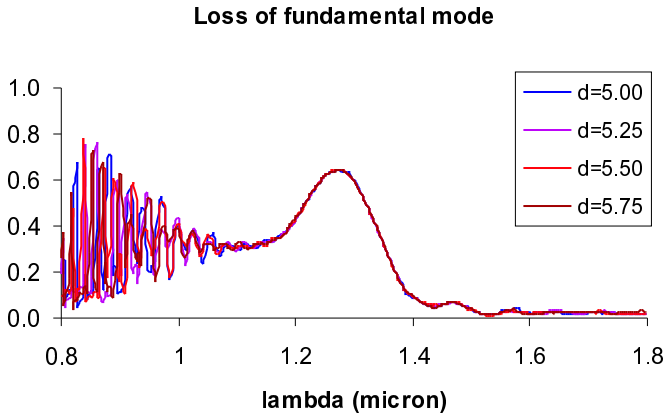

Figure 3.5: Loss of the fundamental mode of the waveguide grating.

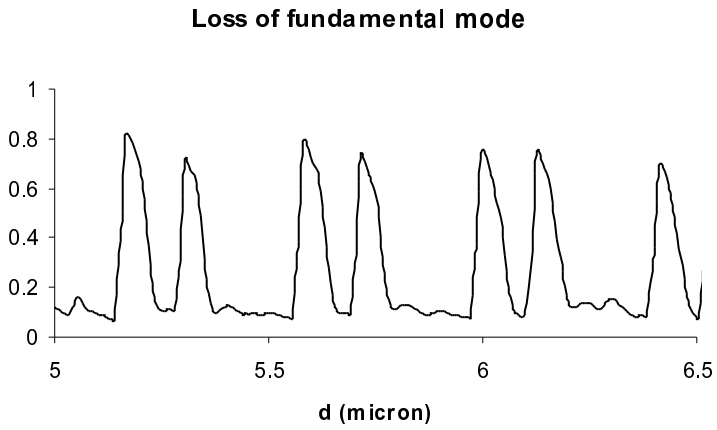

Figure 3.6: Loss of the fundamental mode as a function of the distance to the metal wall.

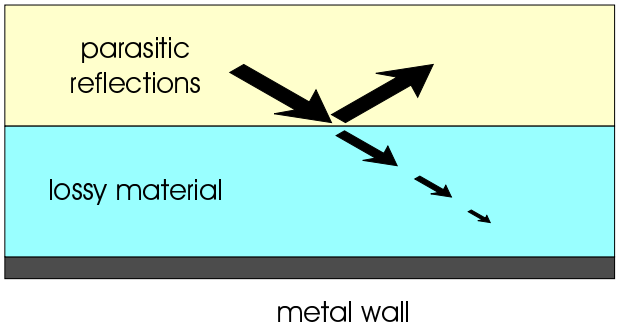

Figure 3.7: Lossy material as an absorber, with parasitic reflections occuring at the interface with the absorber.

An approach that seems attractive at first sight is to place a lossy material in front of the wall. In this way, any radiation that propagates in this lossy medium will be damped, hopefully to such an extent that all the power is absorbed before it has any chance to reflect at the metal wall. The problem here is that now it is possible for the light to reflect at the interface with the absorber, rather than at the metal wall (fig. 3.7). We could tweak the parameters of the lossy medium to minimise this reflection, but this is time-consuming, dependent on the geometry of the structure we want to model and difficult to achieve over a wide range of incidence angles and wavelengths [Noble1998]. In the next section, we will introduce another approach to absorbing radiation that does not suffer from this problem: the use of perfectly matched layers (PMLs).

3.3 Perfectly matched layers

A perfectly matched layer or PML is an artificial material that can absorb radiation without any parasitic reflections at its interface, regardless of wavelength, incidence angle or polarisation. The concept of PMLs can be introduced in a variety of ways. In this work, we will opt for the formalism of complex coordinate stretching, because it allows us to reuse the formulas presented in chapter 2 virtually without any modification. The other ways of describing a PML will be introduced briefly in the next section, together with a historical overview describing the evolution of the use of PML. For more details on these other formalisms, we refer to the Ph.D. dissertation [Derudder2001a].

3.3.1 A brief history of PML

PML was originally introduced in the context of FDTD simulations in [Bérenger1994]. This seminal paper used the so-called 'split field' formalism, in which e.g. the H_z field component is artificially split up into $H_z = H_{zx} + H_{zy}$. This leads to a modified version of Maxwell's equations, where the introduction of split fields provides extra degrees of freedom that can be used to achieve a perfect reflectionless matching at the absorber interface. This produced nothing short of a revolution, since it quickly became apparent that PML vastly outperformed any boundary condition previously known.

In the subsequent years, a flurry of papers was produced on the properties of PMLs. In [Sacks1995], it was shown that the seemingly ad hoc introduced, non-physical, non-Maxwellian concept of split fields was in fact equivalent to a 'regular' physical medium, albeit an anisotropic one with rather peculiar permittivity and permeability tensors. In planar stratified media, these tensors take the form $\varepsilon = \varepsilon_0 \varepsilon_r \Lambda$, $\mu = \mu_0 \mu_r \Lambda$ with

$$\Lambda = \begin{bmatrix} \alpha & & \\ & \alpha & \\ & & \alpha^{-1} \end{bmatrix} \tag{3.6}$$

Here, α is a complex number with a negative imaginary component providing absorption. By explicitly splitting α into its real and imaginary components in eq. 3.6, it can be seen that the z-compoment of the tensor has a positive imaginary part. This means that PML is in fact an active medium containing sources, which unfortunately makes it rather unsuited for actual fabrication.

A very elegant way of looking at PML was introduced in [Chew1994], which states that PML can also be considered as a regular isotropic medium, but this time with a complex thickness. We will discuss this in depth in the next section, but we can already provide a hand-waving

argument that an imaginary thickness is able to damp radiation. Consider the formula describing a plane wave:

$$e^{-jk_0 nz} \tag{3.7}$$

This plane wave can be damped by letting it propagate in a medium with a complex refractive index n, as already mentioned in section 3.2.2. However, since both n and the spatial coordinate z appear on equal footing in eq. 3.7, it is clear that a complex thickness will just as well provide absorption. Moreover, this can be done without any parasitic reflections, as we will explain in section 3.3.2.

All this work focused on applications in FDTD methods. The spatial discretisation used in those methods meant however that the PML layer had to be discretised as well. Going from a continuous, analytic PML to a discrete PML can introduce parasitic reflections that are solely due to the spatial discretisation. Therefore, careful attention had to be paid to the absorption profile $\alpha(z)$ inside the PML [Lazzi1997].

In [Derudder1998], our colleagues from the electromagnetics group at our department showed that PML could also be very useful in eigenmode methods. They modelled discontinuities on microstrip lines, using a stepwise approximation of an anisotropic PML.

We built upon their work and expanded it to circular geometries [Bienstman2001d], in order to model VCSELs, and later also to arbitrary slab waveguides [Bienstman2001c], in order to model photonic crystal devices. We opted for a full analytic treatment involving complex coordinate stretching, rather than for a stepwise anisotropic PML, in order to eliminate the reflections due do the spatial discretisation. This also had the additional advantage of not having to engineer the absorption profile in the PML, as we will see later.

3.3.2 Complex coordinate stretching

Consider the situation depicted in fig 3.8. A structure that we want to study is enclosed by a metal wall. This structure has an outer cladding layer with thickness d_0 and index n. Between the cladding and the wall, we insert the PML, a layer also with refractive index n, but with a complex thickness $d_{1,re} - jd_{1,im}$. We want to prove in this section that a wave propagating in the cladding will be absorbed in the PML, without any reflections occurring at the cladding-PML interface. To this end, we first define a one-to-one mapping of the traditional z variable to a complex stretched $\widetilde{z}$ variable [Chew1997]:

$$\widetilde{z} = \int_0^z s_z(z')dz' \tag{3.8}$$

To make this more concrete, we have drawn in fig. 3.8 the case were $\Re(s_z) = 1$ everywhere, and $\Im(s_z) = -a$ in the PML and zero everywhere

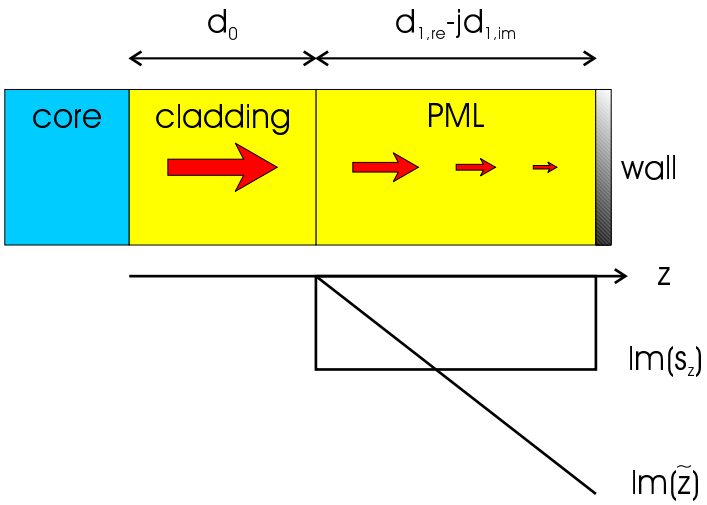

Figure 3.8: Complex coordinate stretching.

else. This leads to $\tilde{z}$ being equal to z, except in the PML, where $\tilde{z}$ has a linearly increasing imaginary component (in absolute value). The use of an integral in eq. 3.8 assures that $\tilde{z}$ varies smoothly, even if s_z has discontinuities.

Now, we take Maxwell's equations and analytically extend them to the complex spatial domain. This means that the equations remain identical, but the coordinates are allowed to assume complex values.

$$\tilde{\nabla} \times \mathbf{E} = -j\omega\mu\mathbf{H} \tag{3.9}$$

$$\tilde{\nabla} \times \mathbf{H} = j\omega\varepsilon\mathbf{E} \tag{3.10}$$

The stretched nabla operator is

$$\tilde{\nabla} = \frac{\partial}{\partial x}\mathbf{u}_x + \frac{\partial}{\partial y}\mathbf{u}_y + \frac{\partial}{\partial \tilde{z}}\mathbf{u}_{\tilde{z}} \tag{3.11}$$

At this point we can already understand why PML provides reflectionless absorption. From a purely formal point of view, the cladding combined with the PML is a *uniform* medium with refractive index n everywhere, but with stretched coordinates. Because coordinate stretching does not formally change the appearance of Maxwell's equations, we can continue to use all the well-known solutions for these equations in unstretched coordinates. For uniform media, we know that these solutions are plane waves, which of course do not reflect when propagating in a uniform medium. At the same time, the imaginary component of the stretched $\tilde{z}$-coordinate absorbs the propagating wave, as we already discussed using eq. 3.7.

We can give this explanation a mathematically more rigorous foundation by deriving the wave equation for the stretched medium. Once again, this is formally identical to the wave equation in unstretched coordinates:

$$\widetilde{\nabla}^2\phi + k_0^2\phi = 0 \tag{3.12}$$

Introducing eq. 3.8 leads to

$$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{1}{s_z}\frac{\partial}{\partial z}\frac{1}{s_z}\frac{\partial}{\partial z} + k_0^2\right)\phi = 0 \tag{3.13}$$

Proposing a solution $\phi = e^{-jk_x x}e^{-jk_y y}\widetilde{\phi}(z)$, this becomes

$$\left(\frac{1}{s_z}\frac{\partial}{\partial z}\frac{1}{s_z}\frac{\partial}{\partial z} + k_z^2\right)\widetilde{\phi} = 0 \tag{3.14}$$

where $k_z^2 = k_0^2 - k_x^2 - k_y^2$. We try a solution of the form

$$\widetilde{\phi}(z) = e^{\pm jf(z)} \tag{3.15}$$

such that $f'(z) = k_z s_z(z)$. Now, we have

$$\frac{1}{s_z}\frac{d}{dz}\widetilde{\phi}(z) = \pm jk_z\widetilde{\phi}(z) \tag{3.16}$$

$$\frac{1}{s_z}\frac{d}{dz}\frac{1}{s_z}\frac{d}{dz}\widetilde{\phi}(z) = -k_z^2\widetilde{\phi}(z) \tag{3.17}$$

which satisfies eq. 3.14. Therefore, the general solution of eq. 3.14 is

$$\widetilde{\phi}(z) = Ae^{-jk_z\int_0^z s_z(z')dz'} + Be^{+jk_z\int_0^z s_z(z')dz'} \tag{3.18}$$

The crucial thing to realise here is that both terms satisfy eq. 3.14 seperately. This means that both the forward and the backward propagating solution are uncoupled. Put differently, one of them propagating will not excite the other one. Therefore, they can propagate without reflection, and this irrespective of wavelength, incidence angle or polarisation.

Finally, one can wonder if this rather ad hoc introduction of materials with complex thickness has any physical bearing whatsoever. In fact, as already mentioned previously, a description in terms of complex coordinates is equivalent to the description of PML as an anisotropic material [Teixeira1997]. This material can of course be described completely by the traditional Maxwell's equations, without resorting to unphysical constructs like stretched coordinates. However, the mathematics of using anisotropic materials are much more involved, which is why we opted for using complex coordinates.

3.3.3 Mode structure

Before illustrating the effectiveness of PML as an absorber, let us first investigate the influence it has on the propagation constants and field profiles of eigenmodes.

The dispersion relation of a slab clad with PML is identical to the one without PML, except for the fact that the outer media have a complex thickness. The dispersion relation can therefore also be cast in the general formalism of eq. 3.3.

Fig. 3.9 shows the distribution of the TE modes in the n_{eff}-plane, for a GaAs waveguide ($n = 3.5$) with a thickness of 1 μm, surrounded on both sides by an air cladding with real thickness of 2 μm. The operating wavelength is 1.55 μm. Three figures are shown, for different imaginary thicknesses of the cladding (-0.2, -0.4 and -0.6 μm). A number of conclusions can be drawn from these pictures. Let us first turn our attention to the guided modes, located on the real axis. Guided modes have exponentially decreasing tails, because they have an imaginary wavevector component in the direction perpendicular to the wall. If the cladding is thick enough, these tails will be negligible by the time they reach the PML, so they should not be influenced by it. We see indeed from the figures that the presence of PML does not change the location of the guided modes. Moreover, it is interesting to point out that the imaginary PML thickness cannot absorb the tails of the guided modes, because of their imaginary wavevector in eq. 3.7. Only the real part of the PML thickness is effective in absorbing these modes.

The opposite is true for the radiation modes, with a real wavevector component perpendicular to the wall. The imaginary thickness absorbs these modes, so that they are no longer located on the coordinate axes, but in the complex plane. In fact, when the absorption in the PML is strong enough, the modes clearly start to cluster in two groups: a first group, where the modes eventually converge to a final location for increasing PML absorption, and a second group, where the modes keep on moving further into the complex plane as the absorption increases. There is a profound reason for this behaviour, but we will have to wait until section 3.4 to fully understand it.

Let us now turn our attention to the field profiles of such modes. Fig. 3.10 plots the absolute value of the **E**-field for a radiation mode with $n_{eff} = 0.416 - 1.600j$ in the case of an imaginary cladding thickness of -0.4 μm. Fixing the total imaginary cladding thickness seems to leave us with the freedom of choosing the exact stretching profile to achieve this total thickness. Therefore, we plot two cases, both with a piecewise constant stretching profile s_z. In the first case, we choose the real thickness of the PML to be half of the cladding layer (1 μm), with the other half remaining unstretched. In the second case, we halve the PML thickness, but at the same time double its absorption. The field profiles in the central region are the same for both cases. In the

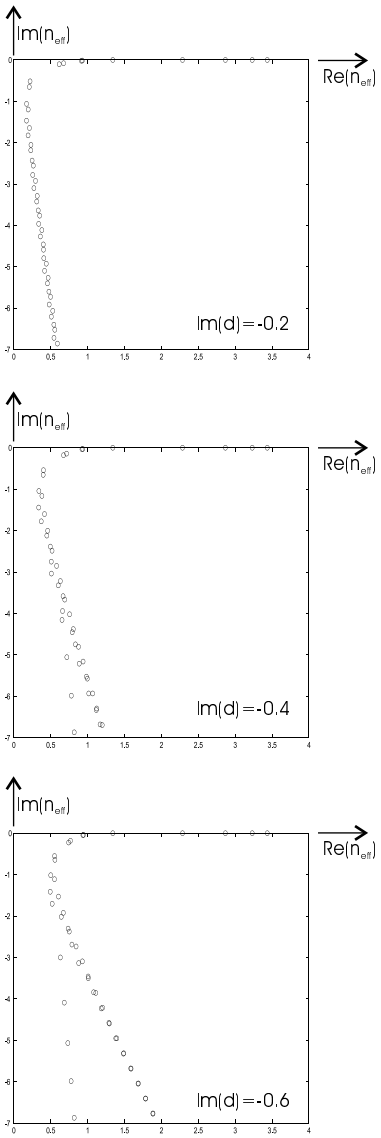

Figure 3.9: Distribution of eigenmodes for different absorption levels in the PML.

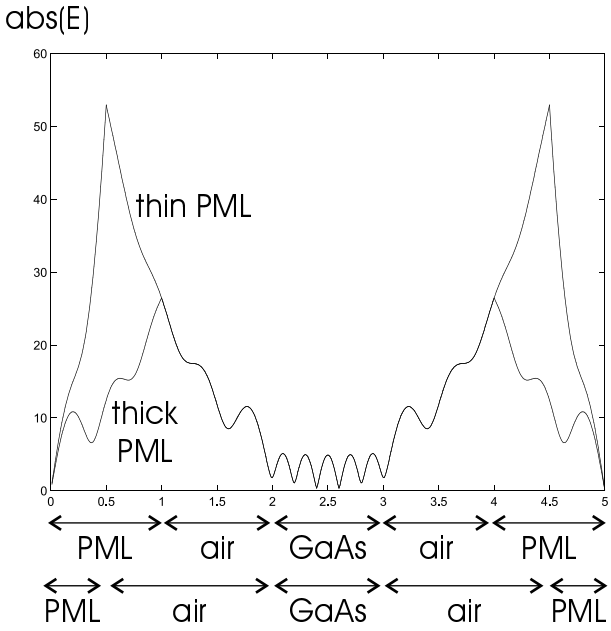

Figure 3.10: Field profile of a radiation mode in the presence of PML.

cladding region, the field increases, but as soon as it enters the PML, it decreases exponentially to fall to zero at the metal wall. For the thin PML, the field is allowed to increase further into the cladding, but this is exactly compensated by a stronger absorption inside the PML.

We want to stress that for the calculation of the dispersion relation, of the overlap integrals and of the scattering matrices, the exact stretching profile is *totally irrelevant*, because all the expressions there involve only the total complex thickness of the uniform medium formed by the combination of cladding and PML. This is a huge advantage compared to the use of PML in e.g. FDTD, where the stretching profile has to be carefully engineered in order to minimise reflections due to the spatial discretisation.

Finally, we want to point out that a similar mode structure exists for TM polarisation and for circular geometries. Fig. 3.11 plots the TM eigenmodes (crosses) alongside with the TE eigenmodes (circles) for the waveguide from fig. 3.9 with an imaginary cladding thickness of -0.4 μm. Once again, there is a branch containing modes that do not move for increasing PML absorption, and this branch is located closer to the imaginary axis than any TE branch. On the other hand, the second TM branch lies further in the complex plane than any TE branch. For the corresponding circular structure (i.e. with a core diameter of 1 μm), TE and TM modes are no longer uncoupled, but combine to form

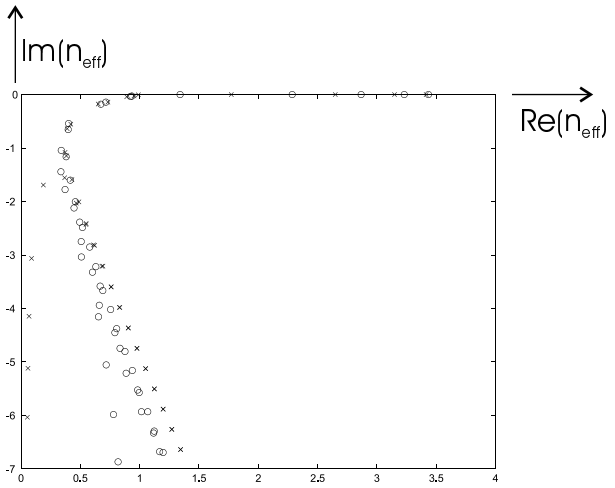

Figure 3.11: Distribution of eigenmodes for TE (circles) and TM (crosses) slab modes.

hybrid modes if the Bessel order n differs from 0. Fig. 3.12 shows the mode structure for $n = 1$. The two branches closest to the imaginary axis are insensitive to further increase in the absorption, the other two are not.

3.3.4 Examples

Let us now return to the example of the waveguide grating, and see if adding an imaginary component to the cladding thickness will solve the problem of parasitic reflections. Fig. 3.13 shows once again the radiation loss of the fundamental mode, but for different values of the imaginary thickness. We see indeed that the PML is successful in damping the oscillations from the parasitic reflections. For a moderately strong PML ($d_{im} = -0.1$ μm), there are still a few residual oscillations left. Although there are no parasitic reflections at the PML interface, the field is not damped sufficiently by the time it reaches the metal wall, so that some part of the radiation will still make it back to the structure under study. In summary, we can say that although PML is perfectly matched (no reflections at its interface), it is not a perfect absorber, because its absorption is still finite[1]. Increasing the imaginary thickness to -0.2 μm is however sufficient to get rid of the remaining oscillations.

To further illustrate the beneficial effects of PML, we plot in fig. 3.14 the field profiles both with and without PML for $\lambda = 0.85$ μm. Without PML, there clearly is a resonance building up in the claddings, due to

[1] We will discuss PMLs with infinite absorption in section 3.4.

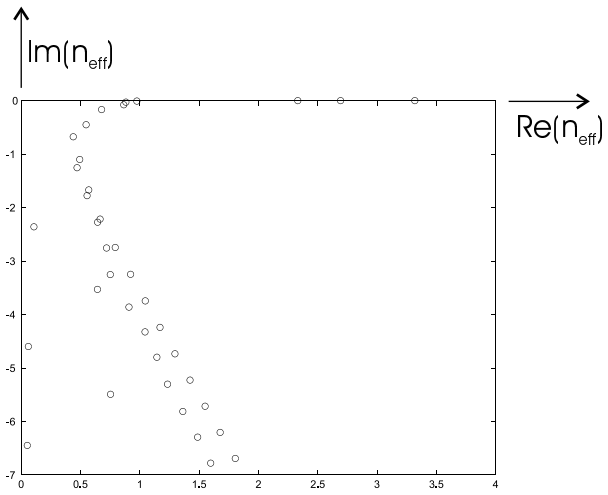

Figure 3.12: Distribution of eigenmodes for a circular waveguide.

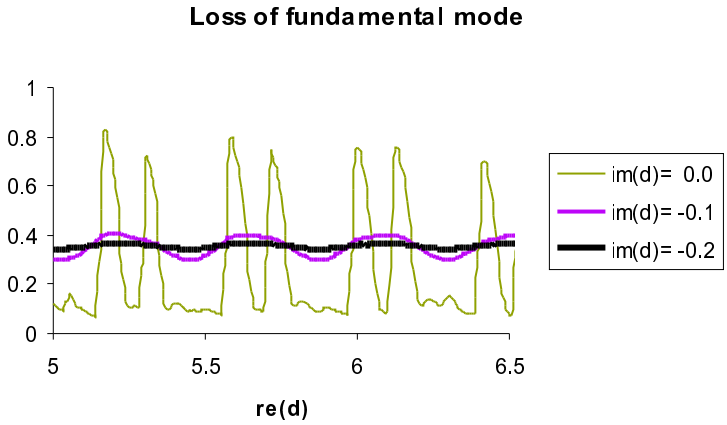

Figure 3.13: Loss of the fundamental mode as a function of the distance to the metal wall, for different imaginary cladding thickness.

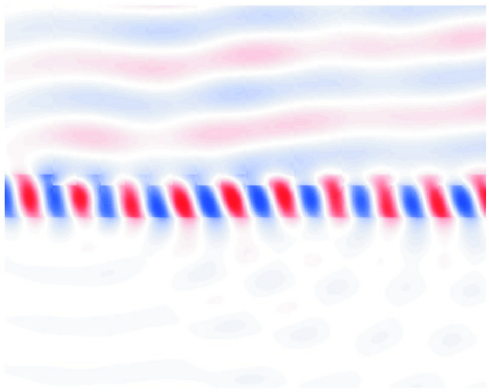

with PML

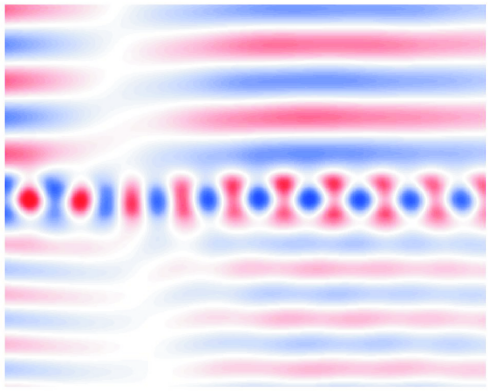

without PML

Figure 3.14: Field profiles for the waveguide grating, both with and without PML.

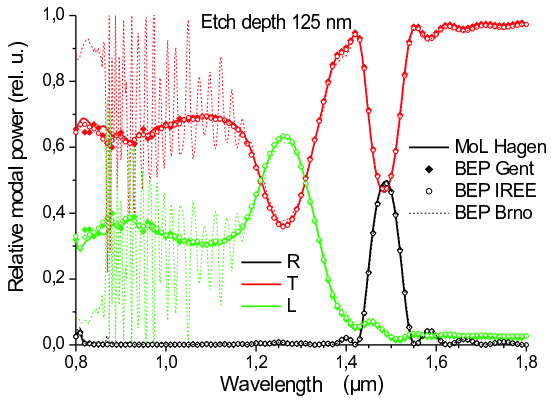

Figure 3.15: Comparison of different models using the waveguide grating structure.

parasitic reflections from the walls. This resonance disappears when PML is introduced. Moreover, when observing the time evolution of these field patterns, it quickly becomes obvious that the waves in the cladding are standing waves without PML and purely travelling waves in the presence of PML.

Finally, since the structure from fig. 3.3 was the subject of a European modelling benchmark, we can validate our results against those of other models. This is illustrated in fig. 3.15. Shown are results using the Method of Lines and several implementations of the eigenmode expansion method (called here BEP for bidirectional eigenmode propagation). The implementation at IREE (the Institute of Radio Engineering at Prague) is modelled after ours and includes PML. The one implemented at the University of Brno does not use PML and clearly suffers from parasitic reflections. Apart from that, all the models are in excellent agreement.

3.4 Leaky modes

Eq. 3.3 also contains the dispersion relation for the open structure, i.e. without any walls. To retrieve this dispersion relation, we simply have to impose that there are no reflected waves in the cladding by putting the r coefficients equal to zero. It then follows immediately that $T_{22} = 0$, which is the required dispersion relation.

We already mentioned in section 2.2.3 that the dispersion relation of the open structure can have solutions in the complex plane. These so-called leaky modes are unphysical, because they increase exponentially towards infinity in the transverse direction. They can be considered as guided modes below cut-off, because they will evolve into proper guided

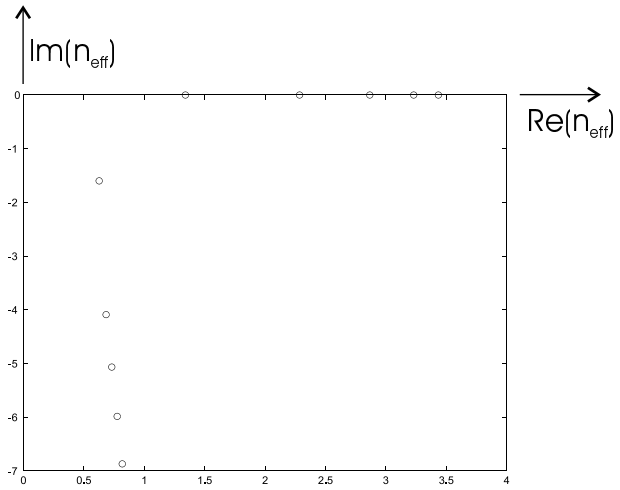

Figure 3.16: Distribution of guided and leaky modes.

modes as the frequency increases. In fig. 3.16, we plot the distribution of these modes in the n_{eff}-plane, for the same waveguide as in fig. 3.9.

Although these modes are non-physical, they form a discrete set and people have tried to use them in eigenmode expansion techniques. The hope here is that every leaky mode can represent a subset of the radiative continuum in open structures. In this section, we will replicate some of the results in literature, where this approach has been successfully applied to model certain structures. We will also investigate whether or not it is possible to extend this technique to a larger class of structures. Finally, we will point out the link connecting leaky modes and PML, leading to a deeper understanding of the properties of modes in the presence of PML.

3.4.1 Normalisation and overlap integrals

When we want to apply the formalism of chapter 2 to leaky modes, we quickly run into a problem concerning the overlap integrals. Indeed, these modes have infinite fields at infinity, meaning that the overlap integrals will diverge. A couple of techniques exist to deal with this issue [Lee1995], the most elegant of which is given in [Sammut1976]. The method presented there consists of deforming the integration path used to calculate the overlap and normalisation integrals. Normally, this path runs along the real x-axis, but it is deformed into the complex plane, as shown in fig. 3.17. The analytical solution of the overlap integrals from the appendices in chapter 2 depends only on the beginning and end points of the interval. Therefore, the precise integration path is not important, only the end points are. If the damping provided

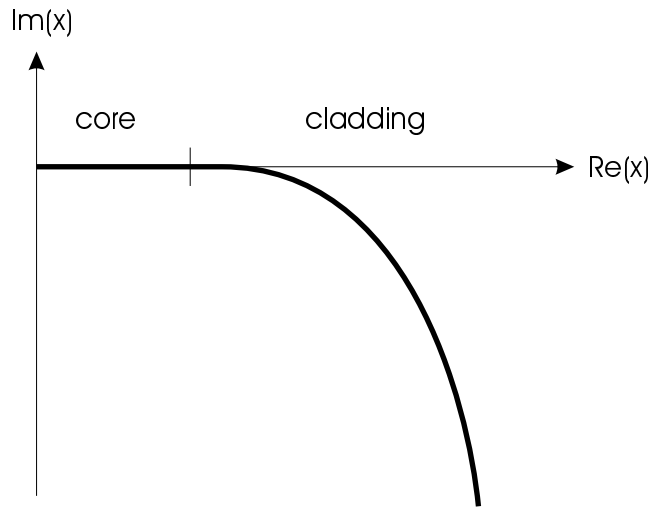

Figure 3.17: Deformation of the integral path to calculate overlap integrals for leaky modes.

by the complex coordinate at infinity outweighs the field increase in the cladding, then this contribution at infinity vanishes and the total overlap integral remains finite.

3.4.2 Example: twin waveguide

We now use leaky modes to model the structure from fig. 3.18 [Lee1995]. It is a twin waveguide section embedded in a single waveguide, which can be used as an interference type device like a directional coupler or a Mach-Zehnder interferometer. All the waveguide sections have $n = 3.381$, while the background has $n = 3.38$. The cores have a width of 4 μm, while the cladding section separating the twin guides has a thickness of 8 μm. The operating wavelength is 1.15 μm. Once again, the propagation is along the z-direction.

Fig 3.19 shows the transmission of the fundamental even guided mode through this structure, as a function of the length L of the central section. A magnetic wall was used to exploit the mirror symmetry along the central axis. One curve is calculated using only a guided mode and a single leaky mode. The other curve uses a cladding thickness of $50 - 0.1j$ μm and 20 modes.

We can compare our curves to the ones published in [Lee1995]. When we first consider their exact results, i.e. those not involving leaky modes, we can conclude that our PML simulations are in excellent agreement. In [Lee1995] however, hard walls were used and the authors claimed they needed many hundreds of modes to achieve con-

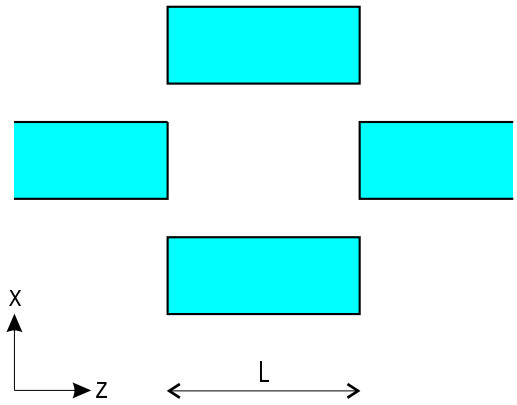

Figure 3.18: Twin waveguide.

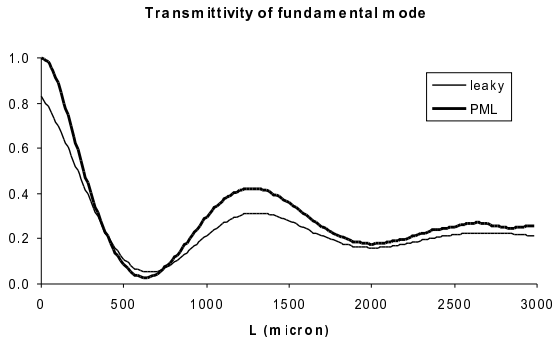

Figure 3.19: Transmittivity in the twin waveguide structure. The PML curve coincides with the exact calculations from [Lee1995].

vergence, which would have resulted in significantly longer calculation times. We on the other hand needed only 20 modes thanks to the use of PML.

If we compare the leaky mode approximations, it turns out that ours performs less well than the one in [Lee1995], but still follows the trend of the exact results quite well. The difference between the two leaky mode approximations are most likely due to the different way of normalising the leaky modes and calculating the overlap integrals. Further research is needed in order to evaluate and compare these techniques more thoroughly.

There are however still a few issues with both leaky mode expansion methods. First of all, for $L = 0$ our approach gives a transmission of 85%, instead of the correct value of 100%. In the original paper [Lee1995], a value of 120% was obtained, equally wrong, but different. We can therefore conclude that leaky mode expansion has some trouble in accurately representing the near field. Secondly, in the cladding far from the waveguides, there still is an unphysical exponential field increase. This means that the leaky mode formalism has a domain of validity limited to regions of space close to the core. This point of view is also expressed in [Snyder1983], where a critical angle θ_c with the z-axis is derived, beyond which the field profiles can no longer be considered as physical.

It has to be stressed that the twin guide structure is relatively easy to model, because of the very weak index constrast that causes little or no backreflection or scattering. The loss mechanism in this device is therefore very similar to that of an infinitely long twin waveguide, and leaky modes seem to be able to capture these physics well. For the grating waveguide structure we presented earlier, we were unable to use this method successfully. There however, the mechanisms at work are completely different. The strong index contrast causes substantial scattering, and the amount of loss is very strongly dependent on the period of the grating, rather than on the natural leakage of long unperturbed waveguide sections. Also, the grating period obviously has near-field effects, which seem to be difficult for the leaky mode expansion to model.

Moreover, it is possible to prove conclusively that the set of leaky modes is not complete (see e.g. [Smith1995]). In view of this, it is quite ill-advised to rely solely on leaky mode expansion without any verification from another method.

3.4.3 Relation to PML

The path deformation used to normalise leaky modes is strongly reminiscent of the complex coordinate stretching PML technique. This suggests a closer relation between the two techniques, which we will try to

uncover in this section.

Let us consider once again the general slab waveguide from fig. 3.1 and its general dispersion relation eq. 3.3, in which we have factored out $e^{+jd_1 k_{x,1}}$ and $e^{+jd_2 k_{x,2}}$. We will assume a very strong PML present in both claddings.

$$\begin{bmatrix} r_N e^{-2jd_N k_{x,N}} & -1 \end{bmatrix} \cdot \begin{bmatrix} T_{11} & T_{12} \\ T_{21} & T_{22} \end{bmatrix} \cdot \begin{bmatrix} r_0 e^{-2jd_1 k_{x,1}} \\ 1 \end{bmatrix} = 0 \tag{3.19}$$

If a certain mode is strongly attenuated in the top PML layer, then $d_N k_{x,1}$has a large negative imaginary component and $e^{-2jd_N k_{x,N}} \ll 1$. Similarly, $e^{-2jd_0 k_{x,0}} \ll 1$ for strong absorption in the bottom PML. We can now distinguish between four cases, depending whether a mode is strongly absorbed in both PML layers, only the top one, only the bottom one, or neither of them.

For the first set of modes, we assume strong absorption in both PMLs, such that most of the field is located in the central regions of the slab:

$$\begin{bmatrix} 0 & -1 \end{bmatrix} \cdot \begin{bmatrix} T_{11} & T_{12} \\ T_{21} & T_{22} \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 1 \end{bmatrix} = 0 \tag{3.20}$$

This is satisfied when $T_{22} = 0$, which is nothing other than the dispersion relation of the open structure. Mathematically, this is also obvious, since infinite PML absorption ($\exp(-2jdk) = 0$) has the same effect in eq. 3.20 as considering a truly open structure ($r = 0$). This is also what one expects from a physical point of view: as long as the absorption in the PML is high enough, no wave will make it back to the core after reflecting at the outer walls. Therefore, viewed from the core, a very strong PML cannot be distinguished from a truly open structure.

We already know what kind of modes are found with the dispersion relation of the open structure: on one hand the guided modes, but on the other hand also the unphysical leaky modes. This is precisely why some modes converge to a fixed location in fig. 3.9: they converge to the leaky modes of the open structure for sufficienty large absorption in the PML.

For the second set of modes, we assume a field strongly damped in the PML at $x = 0$, but not so in the other PML layer. We now get for the dispersion relation

$$T_{12} r_N e^{-2jd_N k_{x,N}} - T_{22} = 0 \tag{3.21}$$

which can be recast as

$$e^{-jd_N k_{x,N}} \cdot r_N \cdot e^{-jd_N k_{x,N}} \cdot \frac{T_{12}}{T_{22}} = 1 \tag{3.22}$$

This equation has a clear physical significance. The first factor describes propagation over a distance d_N in the PML, the second factor is the reflectivity at the wall, the third one describes the wave as it propagates from the wall back to the core, and the final term is the reflection of the core as seen from the top.[2]. In other words, eq. 3.22 clearly imposes a round-trip gain of unity in the cavity formed by the top wall and the central slab.

A similar set of modes can be found that resonate in the bottom PML. Because the properties of these modes obviously depend more strongly on the PML than on the central core region, these modes are sometimes called Bérenger modes [Derudder2001b]. These Bérenger modes are the ones in fig. 3.9 that move further into the complex plane as the PML absorption increases.

Finally, the case where the modes are neither absorbed in the top nor the bottom PML is purely academic, since it contradicts our initial assumption of a strongly absorbing PML.

3.4.4 PML with infinite absorption

In the previous section, we showed that for strong absorption in the PML, some modes converge to the leaky modes of the open structure. This result was first presented in [Rogier2001], but for the less general case of a microstrip substrate.

Now, we want to take the argument one step further, and show that for *infinite* absorption in the PML, eigenmode expansion becomes *completely* equivalent to leaky mode expansion. In order to prove this, we need to elaborate on a couple of points.

First of all, for an infinitely strong PML, the modes from eq. 3.20 coincide exactly with the leaky modes of the open structure.

Secondly, the Bérenger modes vanish entirely if the absorption of the PML becomes infinite. This can be seen from eq. 3.22 and its corresponding physical picture that a resonator with infinite losses cannot sustain any modes. This also follows from the explicit formulas of the Bérenger modes presented in [Rogier2001] for a microstrip substrate, where no non-trivial solutions remain for infinite absorption.

Lastly, for infinitely absorbing PML, the calculation of overlap integrals proceeds in exactly the same way as in the case of leaky modes. Indeed, both use complex coordinate stretching to achieve absorption. For an infinitely strong PML, the contribution to the overlap integrals at the walls vanishes, just as the contribution at infinity in the case of leaky modes.

Together these arguments show that leaky mode expansion is nothing else than using a PML with infinite absorption. This sheds some

[2] This can be seen immediately from $\begin{bmatrix} r \\ 1 \end{bmatrix} = \begin{bmatrix} T_{11} & T_{12} \\ T_{21} & T_{22} \end{bmatrix} \cdot \begin{bmatrix} 0 \\ t \end{bmatrix}$

more light on the physics behind each of these methods.

This also leads to some important observations with respect to the completeness of the set of eigenmodes in the presence of PML. If we assume an otherwise lossless structure, increasing the absorption in the PML leads from a set that is proven complete (for zero absorption) to a set that is proven incomplete (for infinite absorption). This suggests that while the set of modes in the presence of PML might not be mathematically complete, for all intents and purposes it is still close enough to complete if we restrict ourselves to moderate absorption. All the empirical evidence of the simulations performed in this work seems to show that this is indeed the case: for moderately strong absorption, using PML modes seems to be a very valid and practical approach. By moderate absorption, we mean a situation whereby the modes are still located relatively close to the coordinate axes, and where there is not yet a clear distinction between the Bérenger and the leaky branch.

Finally, it is worth noting that, just like leaky modes, PML modes exhibit an exponential field increase in the cladding before entering the absorbing region (see e.g. fig. 3.10). This also indicates that the validity of the fields in the presence of PML seems to be limited to an area of the cladding immediately around the central region.

3.5 Transparent boundary conditions

3.5.1 Derivation using the orthogonality requirements

Now would be a good time to revisit the eigenmode orthogonality requirements at the boundaries (eq. 2.15):

$$\int\int_{\delta S} (\mathbf{E}_1 \times \mathbf{H}_2 - \mathbf{E}_2 \times \mathbf{H}_1) \cdot d\mathbf{S} = 0 \tag{3.23}$$

For the TE modes of the general slab waveguide and at the lower boundary at $x = 0$, this becomes

$$\int (E_{y,1} H_{z,2} - E_{y,2} H_{z,1})\, dz = 0 \tag{3.24}$$

At PEC or PMC walls, this is always fulfilled, since these boundaries impose the vanishing of the tangential electric field E_y or the tangential magnetic field H_z respectively. The presence of PML does not change this conclusion, since the PML is backed by a PEC or a PMC wall that imposes the same boundary conditions. Likewise, the orthogonality holds for the leaky modes as the special case of PML with infinite absorption.

Let us now decompose the fields at $x = 0$ in forward and backward propagating plane waves (see eq. 3.4 and 3.5):

$$E_y = F + B \tag{3.25}$$

$$H_z = \frac{k_x}{k_0 c \mu}(F - B) \tag{3.26}$$

We can then write eq. 3.24 as

$$(F_1 + B_1)\, k_{x,2}\, (F_2 - B_2) - (F_2 + B_2)\, k_{x,1}\, (F_1 - B_1) = 0 \tag{3.27}$$

or as

$$k_{x,2}\frac{F_2 - B_2}{F_2 + B_2} = k_{x,1}\frac{F_1 - B_1}{F_1 + B_1} \tag{3.28}$$

This suggests using the following form of boundary condition, with $k_{x,0}$ an arbitrary constant:

$$k_x \frac{F - B}{F + B} = -k_{x,0} \tag{3.29}$$

Rewriting this in the canonical form $F - rB = 0$ (eq. 3.1) gives us for the reflection coefficient

$$r = \frac{k_x - k_{x,0}}{k_x + k_{x,0}} \tag{3.30}$$

Note that this is a case where the reflection coefficient depends on the incident angle. The absolute value of r is always smaller than or equal to 1, and is identically zero for the angle corresponding to $k_x = k_{x,0}$.

Another way of looking at this boundary condition is to reformulate eq. 3.29 in terms of $E_y = F + B$ and its derivative $E_y^{'} = -jk_x(F - B)$:

$$-jk_{x,0}E_y + E_y^{'} = 0 \tag{3.31}$$

This is a mixed von Neumann-Dirichlet boundary condition with constant coefficients. For $k_x = k_{x,0}$ it is satisfied by $e^{+jk_x x}$, which is a plane wave propagating downward from the core to the wall, without the presence of any reflected wave $e^{-jk_x x}$.

This kind of boundary condition was presented in [Hadley1991], where it was called "transparent boundary condition" (TBC) and applied to the beam propagation method. A similar boundary condition was already presented in a different context in [Leontovich1948]. The TBC is also quite popular among MoL practitioners [Gerdes1992]. Its use in eigenmode expansion methods has to our knowledge never been thoroughly studied, although the technique has been hinted at in [Smith1995].

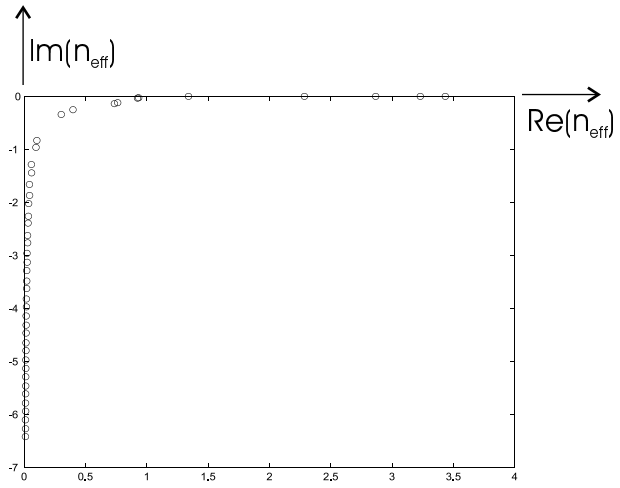

Figure 3.20: Distribution of modes in the presence of a transparent boundary condition.

Fig. 3.20 plots the distribution of eigenmodes for the slab waveguide from fig. 3.9. The cladding is once again 2 μm thick, and the walls are TBCs with a $k_{x,0}$ corresponding to 45 degrees incidence angle.

A striking feature of this figure is that the modes are located much closer to the coordinate axes, compared to e.g. PML and leaky modes.

3.5.2 Example: laser facet

As an example, we take another benchmark problem that has been extensively studied in numerous papers in the literature ([Ikegami1972], [Herzinger1993] and [Haes1996] among others). Its aim is to calculate the reflectivity of the cleaved end facet of a semiconductor laser, with the geometry of fig. 3.21. The core has a width of 275 nm, and we consider the incidence of the fundamental TM mode at a wavelength of 860 nm, propagating horizontally to the facet on the right. An electric wall was used to exploit the symmetry.

In fig. 3.22, we plot the convergence of the reflectivity as a function of the number of modes, and this for different boundary conditions. We compare a PEC wall placed at 2 μm from the core, a PML layer with the same real thickness but with an imaginary thickness of -0.2 μm, and finally a TBC with zero reflection for 45 degrees placed at the same real distance. All curves show a similar behaviour, i.e. they require about the same number of modes to converge to their final value. The reflectivity of the PML and TBC curve is in excellent agreement. The PEC curve however converges to a wrong value because of parasitic reflections. We can also obtain the correct result with a PEC, but only

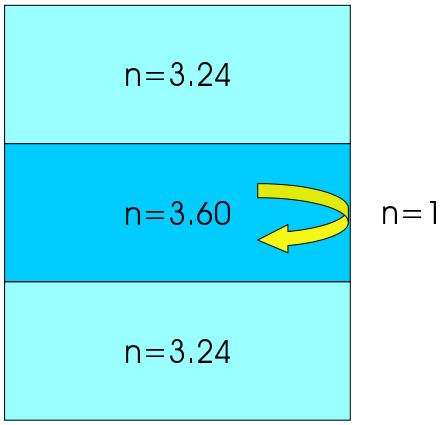

Figure 3.21: Laser facet.

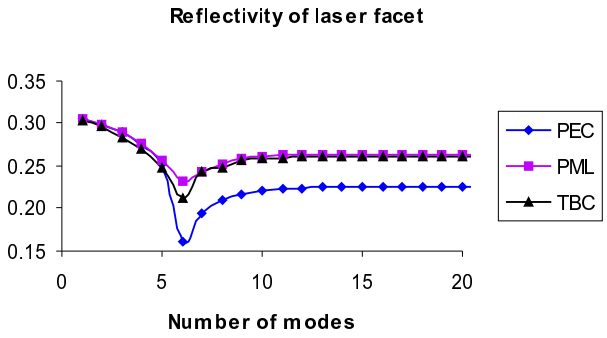

Figure 3.22: Reflectivity of fundamental TM mode at laser facet.

if we place the wall much further from the structure. This is of course at the expense of having to retain a much larger number of modes because of the larger computational domain. Advanced boundary conditions like PML and TBC therefore have an additional benefit next to their increased accuracy: because the walls can be placed much closer to the structure under study, the calculation times can be drastically reduced (remember that these scale as N^3).

3.6 Comparing the different boundary conditions

We will now briefly compare the different boundary conditions presented so far. It is clear that hard walls are not very effective in modelling open space. Their only valid use is the exploiting of symmetries. Leaky mode expansion can give very good approximations with only a small number of modes, but unfortunately this method is only useful for long waveguide sections with small index contrasts. PML boundary conditions are very efficient, and have a much wider domain of applicability. They are reflectionless for any wavelength and provide moderate to strong absorption, making them very suited to model open space. Transparent boundary conditions exhibit zero reflectivities for a single angle and moderate to high reflectivities for other incidence angles. The fact that the zero-reflection angle has to be known in advance is sometimes a disadvantage.

In fig. 3.23, we plot the power reflectivity of some of these boundary conditions: two PMLs with imaginary thickness of -0.2 and -0.3 μm respectively backed by a PEC, and two TBCs with zero reflection at either 30 or 60 degrees. The wavelength considered is 1 μm. The incident fields are plane waves that are propagative in the direction normal to the walls.

From fig. 3.23, it can be seen that PMLs exhibit their lowest reflectivity for normal incidence, because this corresponds to the highest value of k_x in eq. 3.7. A larger imaginary thickness obviously means lower residual reflections. It is clearly visible that the TBCs have zero reflection for a single angle. If this zero-reflection angle moves to larger angles (i.e. towards grazing incidence), their reflectivity for the other angles increases. All the curves share the fact that their reflection approaches unity for grazing incidence. This means that it will still be very difficult to absorb modes close to cut-off.

Finally, it has to be noted that nothing forbids us to combine both PML and TBC in a single boundary condition, in which case the total absorption in dB is the sum of the absorption of each of the individual boundary conditions. Such a combination is probably overkill in most situations, but could be useful for some simulations where extremely

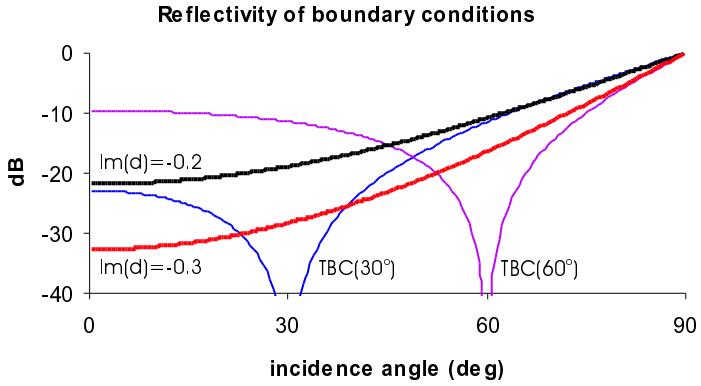

Figure 3.23: Power reflectivity of different boundary conditions.

powerful absorbers are required.

3.7 Locating complex propagation constants

So far, we have not yet elaborated on the method used to find the zeros of the dispersion relation in the complex plane. This is nevertheless very important, because if locating these zeros is too time-consuming, the benefit of using advanced boundary conditions like PML or TBC is essentially lost.

A popular method of complex zero finding is the argument principal method (APM), which was first presented in [Delves1967]. This method numerically calculates contour integrals of the form

$$R_k = \frac{1}{2\pi j} \oint_C z^k \frac{f'(z)}{f(z)} dz \tag{3.32}$$

The coefficients R_k can then be used to form a polynomial which has the same zeros as f within the contour C. Alternatively, one can use the fact that in the absence of poles, R_0 equals the total number of zeros within C. After successive refinements of C to find subregions that contain only a single zero, an iterative method (like e.g. the Muller method [Muller1956]) can be used to locate that single zero [Derudder2001a]. The routines we used that implemented this method were due to [Blomme1996], but they were not fully optimised and not 100% robust in all corner cases, especially when zeroes were located close to the contours.

Therefore, we explored a different technique to locate the modes in the presence of PML [Bienstman2001d]. Without PML, the eigenmodes are located on the coordinate axes[3], where they are readily found using

[3]Apart from the rare complex modes, which have to be dealt with seperately.

e.g. Brent's method [Press1993]. If we now gradually increase the absorption in the PML to the desired value, we can track the modes as they move into the complex plane. This method of root tracking proved to be quite fast and robust, and has been our preferred method to locate PML modes.

Finally, it has to be mentioned that there exist other methods for complex root finding, which seem very promising to investigate and to implement in a future version of the software. The so-called ADR method, adapted to waveguide problems in [Anemogiannis1994], uses integrals of the form

$$G_k = \frac{1}{2\pi j} \oint_C \frac{z^k}{f(z)} dz \tag{3.33}$$

The main advantage of this method is that it does not require the numerical evaluation of the derivative f', making it significantly faster. Moreover, it can be used for contours that enclose poles or branch point singularities, which should make it far more robust that the APM.

Chapter 4

CAMFR

> "To err is human, to really foul things up requires a computer."
> Farmers' Almanac, 1978.

4.1 Introduction

In this chapter we will pay some attention to the software engineering principles that underly the design of our eigenmode expansion CAD tool. These aspects are often overlooked, but are nevertheless of paramount importance if we want to produce modelling tools that are performant, flexible, and easy to maintain and to extend.

The approach we took was to try and produce a framework where different building blocks could be recombined at will, preferably even in ways that were not envisaged originally. It is important to stress that we did not want to focus the software on a single application, but rather conceive it as a flexible framework. This is also reflected in its name CAMFR, short for CAvity Modelling FRamework. The validity of this framework approach will also be illustrated by the next three chapters, where we will combine the CAMFR components to model three different application domains.

In order to achieve such a component level flexibility, it is really helpful to make use of the object-oriented (OO) capabilities of an advanced programming language like C++ [Stroustrup1997]. In the remainder of this chapter, we will discuss some aspects of the object-oriented paradigm that can aid us in achieving these goals. We present this material without any claim to originality or completeness, but merely as an illustration that recently developed software engineering techniques can be applied successfully to the field of scientific computing. A considerable fraction of the members of that community are still relatively conservative and slow to embrace new techniques. Many of

them keep on developing new applications in Fortran77, blissfully unaware of the advances in software engineering in the last decades. It is our hope that this chapter might persuade them to at least give these new developments some consideration.

4.2 Abstract data types

On many occasions, it can be useful for a programmer to define new (abstract) data types and to deal with them in a way that is just as natural as when using the built-in types of a language, like e.g. floating point numbers or integers.

A typical example of this is performing elementary operations on matrices, where one would prefer to write

```
A = B + 2*C
```

rather than something along the lines of

```
matrix_scalar_mult(C, 2, tmp)
matrix_add(B, tmp, A)
```

For relatively simple expressions like these, the difference is still rather small, but for more complicated formulas, it quickly becomes quite difficult to construct and to debug expressions using functions like matrix_add() and their relatives. However, in languages like C or Java, this approach is the only option. C++ on the other hand allows the programmer to redefine operators like + and * and to give them the desired behaviour for the data types under consideration. It is clear that such functionality greatly benefits the clarity of the code, making it easier to develop, read and debug.

Other data structures where this is useful are complex numbers, spatial coordinates or data structures representing the electromagnetic field. Another perhaps unexpected application lies in the definition of waveguide structures. In CAMFR, we can write

```
DBR = air(0) + 20*(GaAs(.070) + AlAs(.084)) + air(0)
```

This definition does exactly what one thinks it does: it describes a 20 period GaAs/AlAs DBR sandwiched between two air layers. The thickness of the layers is 70 nm for GaAs and 84 nm for AlAs. This way of defining structures is very concise and yet very clear, especially when compared to a more traditional and popular approach, where this structure would be defined in a 42 line text file with each individual layer specified on a separate line.

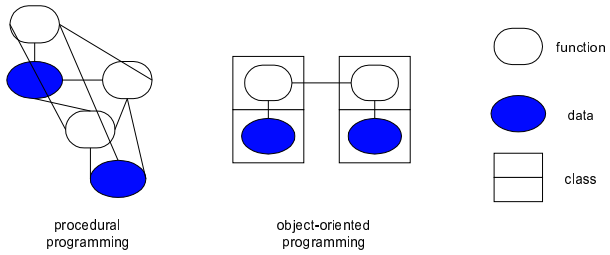

Figure 4.1: Procedural versus object-oriented programming.

4.3 Encapsulation and data hiding

It could be argued that the material presented in the previous section, however useful it may be, is nothing more than syntactic sugar. Object orientation however goes much deeper, and presents a fundamental paradigm shift from the traditional style of procedural programming.

In procedural programming, code is organised as a set of procedures or functions that call each other and that operate on data (fig. 4.1). This application data is often visible throughout the entire code and can be modified at will from within different functions. If not applied carefully, such an approach can quickly lead to an unintelligible mess, when the responsibility for certain data is scattered throughout many functions. Some functions might modify this data in ways that are not expected by other functions and these unanticipated side effects can make the code hard to understand or introduce subtle bugs that are tough to track down. Also, in this approach it is not uncommon that all the functions have to concern themselves with fairly low-level details, such as the precise way in which the data is stored in memory.

In object orientation, the code is constructed as a set of objects that interact in precise and well-defined ways (fig. 4.1). Objects can be seen as abstractions of real-life entities. They exhibit a certain *behaviour*, i.e. they can provide certain services to other objects. These services are expressed as a set of functions (or 'methods' in OO parlance) that the outside world can call on this object. To do their work, objects also have some internal *data*, but this data should not be made accessible directly to other objects. Rather, they should be given the information they need through methods, as part of the functionality provided by the object that owns the data. In this way, the object's data is encapsulated and hidden from the outside world, which prevents the kind of direct tampering that is typical for procedural programming.

To make this less abstract, in CAMFR e.g. we have objects that describe Waveguides and other objects that model the Interface between two Waveguides. The Interface object will obviously need to know about the modal propagation constants in the Waveguides, and will kindly

ask the Waveguides this information by calling the appropriate function. The Interface object should not need to concern itself with the internal details of how the propagation constants are stored inside the Waveguide object. It shouldn't care less if the data is read from disk, from memory, or calculated on the fly. Conversely, a Waveguide should not need to know about the internal details of an Interface, and should continue to work equally well if Interface suddenly decides to use a different algorithm to calculate its scattering matrix. This approach makes it a lot easier and transparent to implement certain optimisations. For example, Interfaces between two uniform Waveguides have diagonal scattering matrices. Interface can then use more efficient algorithms and storage schemes, without us having to change anything to the other classes that use Interface.

In summary, the OO approach strives for a clear decoupling of the 'what' from the 'how'. Other objects expect certain information or behaviour from an object ('what'), but exactly 'how' this functionality is realised is irrelevant for them. The software should be constructed as a set of objects, each of them with a clearly defined responsibility. They do not interfere in the internal workings of other objects ('separation of concerns'). These objects should be designed such that they interact as little as possible, but when they interact, they should do so in well-defined ways, without any unexpected side effects.

These observations might seem trivial, but many programs continue to be written that do not follow these principles. They scatter the responsibility for functionality over a large set of unrelated functions that can interact in complicated ways. For very large projects, this can result in a maintenance and debugging nightmare.

4.4 Polymorphism

Let us return to our example of Waveguides and Interfaces. Obviously, we will want to model different kinds of waveguides, like SlabWaveguide or CircWaveguide, that each take care of the details specific to a given geometry. However, we know from the treatment in chapter 2 that the Interface's mode matching algorithm can be formulated without any reference to the geometry. As long as Waveguides can give Interface the information concerning the overlap integrals, Interface can do its job and does not have to know anything about the underlying geometry.

A language like C++ offers us the tools to express this kind of behaviour. We can define an abstract class Waveguide, that outlines all the functionality the outside world can expect from it, like getting information on propagation constants or on overlap integrals. This abstract class only concerns itself with the 'what' question, it does not provide any implementation details. These details (the 'how') are filled in by SlabWaveguide and CircWaveguide, according to their geometry. Slab-

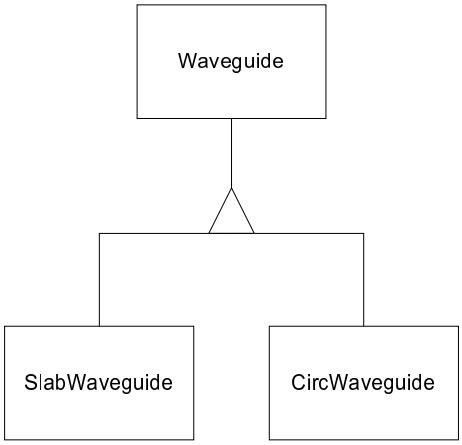

Figure 4.2: The abstract class Waveguide.

Waveguide and CircWaveguide are said to *inherit* from Waveguide, i.e. they are a specialisation of a more general class Waveguide (fig. 4.2).

Inheritance is a very powerful mechanism to express generalisation-specialisation relations, but it is especially useful in combination with a feature called *polymorphism*. This allows an abstract class like Waveguide to dynamically take the form of a SlabWaveguide or a CircWaveguide at run time. This means that we can implement the mode matching in Interface entirely in terms of abstract Waveguides. The class Interface does not even have to know that there exist classes like SlabWaveguide. During the execution of the program, the C++ run time environment will make sure that the function calls on the abstract Waveguide will get passed to the correct concrete class.

The main benefit in this is a substantial amount of code reuse. If we decide in the future to implement an EllipticWaveguide, also derived from Waveguide, we do not have to change *anything* in the implementation of Interface. As far as Interface is concerned, EllipticWaveguide is just another Waveguide, just like SlabWaveguide or CircWaveguide. This makes the Interface class extremely future-proof.

This separation between algorithms that are geometry-dependent and those that are geometry-independent is pervasive throughout our software, and is precisely what gives CAMFR its framework properties.

Another abstract class is that of the Scatterer, which is defined as anything that has a scattering matrix, regardless of its geometry. An Interface is actually also a Scatterer, just like the class Stack, which describes a sequence of Waveguides separated by Interfaces.

The abstract class Scatterer is then used by the extensions that we will describe in the next three chapters: calculating band structures of an infinite periodic repetition of a Scatterer, finding the lasing modes of

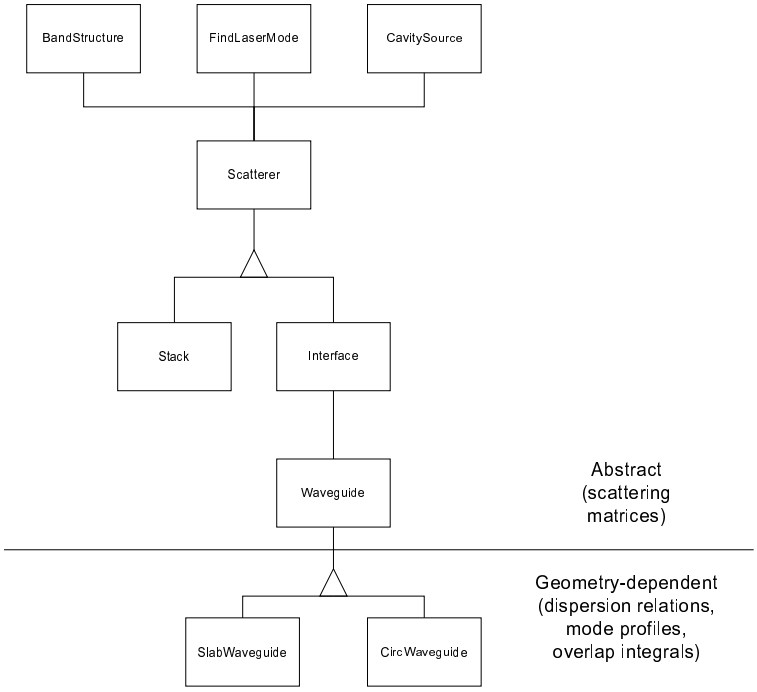

Figure 4.3: Concept diagram for CAMFR.

a Cavity between two Scatterers, or calculating the response of a Cavity to an arbitrary Source. Once again, we stress that these routines only have to be implemented in terms of abstract scatterers, without any reference to the underlying geometry, being it a currently implemented geometry or one that was not envisaged at the time of the original design. This is a significant advantage over a non-OO paradigm, where e.g. the effort of implementing a band structure routine would have to be duplicated N times to write a slightly different routine for each of the geometries.

Fig. 4.3 gives a general overview of how different classes and aspects of CAMFR interrelate. It also shows the separation between geometry-dependent and abstract entities.

4.5 Abstraction penalty

An often heard criticism about the object-oriented approach is that it can lead to slower programs, because the abstraction mechanisms induce a certain amount of overhead (the so-called abstraction penalty). However, there are a number of reasons why this is not as problematic

as one would first think:

- As compiler technology continues to evolve, we can expect the influence of the abstraction penalty to reduce. We already experienced this in practice, where new compiler releases improved the performance of CAMFR without us having to do any recoding.
- Run time is not the only relevant figure of merit in software engineering. Other aspects to take into account are ease of development, whether or not the code is hard to understand and to debug, and if it accommodates future extensions and maintenance without much problems. Object-orientation can be a clear win with regard to these issues.

In conclusion, one could say that in spite of the abstraction penalty, OO is definitely the most future-proof approach when developing software. We can contrast this to the other end of the spectrum, where we would write software in a low-level language like assembler. This would most definitely result in improved performance, but developing the software would take exponentially more time. Moreover, this work would only have a limited useful lifespan, as future versions of the processor or entirely new chip architectures would render it obsolete.

4.6 Computational steering

We already mentioned that CAMFR is constructed as a set of communicating objects or components. However, to run a simulation, it must be specified how these separate pieces should work together in order to obtain results.

Originally, this was done by writing a separate C++ file for each simulation, a file that 'glued' all these components together to describe the structure to be modelled and the calculations to be performed. While this approach worked, it still presented a number of problems. The biggest drawback was that each C++ file specifying a particular simulation had to be compiled before it could be used. This could take a long time, such that for small simulations the compile time could be longer than the actual run time. Clearly, a faster and more flexible solution was called for.

A very popular approach in these kinds of simulations is to describe the structures to be modelled in textual input files that are read in and processed by the main program. These input files usually take the form of a long series of numbers, without any semantics attached to them. This means that accidentally omitting a number or changing the position of two numbers can lead to hard-to-track-down problems, if the software desperately tries to analyse materials with a refractive index

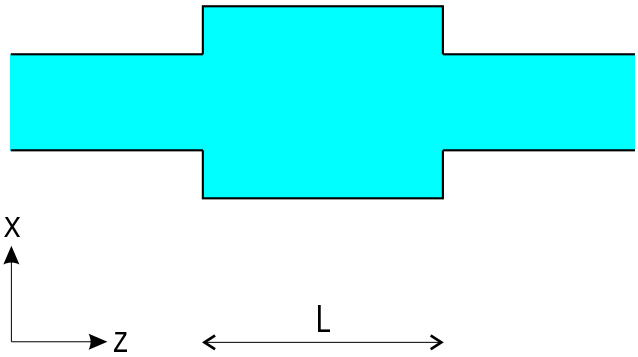

Figure 4.4: Waveguide with a widening and narrowing core width.

of 500 nm. Additionally, when we want to sweep a certain parameter of the simulation, separate input files have to be created for each parameter value, a process that is tedious and error-prone.

What is really needed is an interpreted scripting language (as opposed to a compiled one) that has all the expressiveness of C++ to create e.g. loops and to use compact expressions to define structures as in section 4.2. We could design and implement such a scripting language from scratch, custom-made for CAMFR, but this would be time-consuming and beyond the true focus of our work. Rather, true to the principle that programming is judiciously applied laziness, we chose to integrate CAMFR in the existing scripting language Python [van Rossum2000]. This scripting language is inherently based on object-oriented principles, which makes that it can be seamlessly integrated with an OO framework like CAMFR. Python also boasts a very clean syntax and a gentle learning curve. Moreover, many other pieces of software besides CAMFR have been wrapped in Python, which makes it possible e.g. to perform a simulation in CAMFR and send the results to Matlab for visualisation, all from within the same script file. In other words, Python is acting as a glue language that 'steers' the C++ CAMFR framework and other software packages, and that can even do so in an interactive way.

The Python approach proved to be very valuable to increase the flexibility and improve the turn-around time and visual feedback from the simulations. As an illustration, we present here a Python script to calculate the scattering caused by an abrupt widening and narrowing of the width of a GaAs waveguide (fig. 4.4). The reflectivity of the fundamental mode back to itself is calculated as a function of the length of the wider section. The results are saved to a file and imported in Matlab to visualise them.

```
############################################################
#
# Calculate reflection caused by embedding a section with
# a 0.7 um wide core in a waveguide with a core width
# of 0.5 um.
# Do this for different lengths of this section, write the
# results to a file and plot them in Matlab.
#
############################################################

from camfr import *

# Set constants.

set_lambda(1.5)       # Set wavelength.
set_N(20)             # Set number of modes used.
set_polarisation(TE)  # Set polarisation.

# Define materials.

GaAs = Material(3.5)
air  = Material(1.0)

# Define waveguide sections.

normal = Slab(air(2.0 - 0.1j) + GaAs(0.5) + air(2.0 - 0.1j))
thick  = Slab(air(1.9 - 0.1j) + GaAs(0.7) + air(1.9 - 0.1j))

# Calculate reflection of the fundamental mode for different
# lengths of the thick section.

outfile = open("outfile",'w')

L = 0.0
while L <= .500:
    stack = Stack(normal(0) + thick(L) + normal(0))
    stack.calcRT()
    print >> outfile, L, abs(stack.R12(1,1))
    L += .010

outfile.close()

# Start Matlab and plot the results.

from matlab import *

m = matlab()
m("load 'outfile'")
m("plot (outfile(:,1),outfile(:,2))")
```

4.7 Three-tier architecture

The three-tier or three-layer architecture is a logical extension of the information-hiding principle and the separation of interface from implementation. It identifies three commonly used layers in software, and suggests that they should be decoupled as much as possible. They should be designed and implemented in isolation, without relying on any assumptions concerning the other layers. This facilitates modularity and code reuse.

The three layers in question are:

- *User interface.* This can be a graphical user interface, a scripting language, or any other interface.
- *Core logic.* In the case of CAMFR, all the algorithms implementing the eigenmode expansion model.

- *Low-level numerical routines*[1]. For CAMFR, this is e.g. the LAPACK library that performs linear algebra calculations.

These layers should be orthogonal to each other as much as possible. This means that changing the user interface from a C++ based approach to Python scripts can happen with zero effect on the rest of the code. Also, the core logic should be unaware of the precise nature of the low-level numerical routines. This is accomplished by creating wrapper functions around these routines and by having the core logic call these wrapper functions rather than the underlying numerical code directly. For legacy Fortran numerical routines, this also has the additional benefit that they can be wrapped in C++ functions with a less arcane systax than the underlying original Fortran routine. If we decide later on e.g. to use linear algebra routines optimised for multiprocessor systems, we only have to change the wrapper code, which is located in a single well-defined source file. We do not have to modify the core logic, where calls to the numerical routines are typically scattered throughout many different source files.

4.8 Final remarks

Finally, we want to stress that using an OO-language like C++ is no guarantee for obtaining 'good' software. OO is no 'silver bullet' (see e.g. [Brooks1987]) that instantly solves all software engineering problems. It is perfectly possible to write 'bad' code in C++, just as one can also write 'good' code in C. A shrink-wrapped set of methodologies can never be a substitute for hard work and experience.

More specifically with respect to scientific computing, there is an entire area of expertise that we can only briefly touch upon in this work, and that is the field of numerics and numerical stability. For example, a naive and straightforward implementation of the dispersion relation given by eq. 2.92 will soon turn out to be useless due to numerical instabilities. Identifying and correcting these can be a time-consuming process and requires special attention and care. This kind of work is nevertheless crucial for the success of a numerical model.

[1] For business software, this last layer is instead the storage layer, that deals with how the relevant information is stored in a database system.

Chapter 5

Photonic Crystals

> "The purpose of computing is insight, not numbers."
> R.W. Hamming

In this chapter, we will use the basic building blocks provided by our tool CAMFR to model 2D photonic crystal devices. First of all, we will show in section 5.1 how to calculate the band diagrams of 1D periodic structures. We will also discuss how these techniques can be applied to 2D crystals. In section 5.2, we will introduce line defects in 2D crystals, and calculate the properties of the waveguides formed in this way. We will show how these waveguides can be used to create very compact bends and splitters. Section 5.3 will focus on semi-infinite crystals and will investigate some of the misconceptions in literature regarding the reflection characteristics of these crystals. To end this chapter, we will briefly mention in section 5.4 some of the CAMFR simulations performed by other people in our department on various photonic crystal devices.

5.1 Band structures

5.1.1 Calculating 1D band structures

Consider the structure from fig. 5.1, which consists of an infinite number of repetitions of the same basic period in the z-direction. The structure itself can be two-dimensional, but for the moment we restrict ourselves to 1D periodicity. We know from Bloch theory [Bloch1928] that the solutions of the wave equations in such a periodic medium take the following form:

$$\psi_{k_z}(z) = e^{-jk_z z} \cdot u_{k_z}(z) \tag{5.1}$$

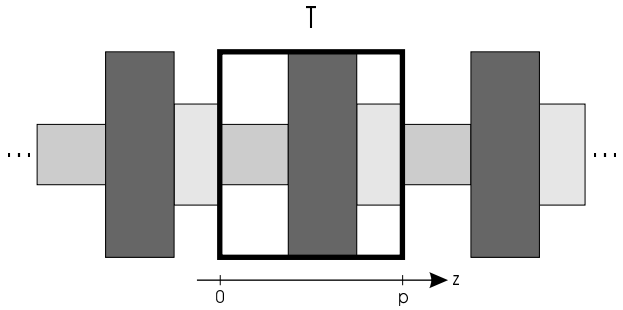

Figure 5.1: 1D periodic medium.

Here, ψ represents any component of the electromagnetic field and u_{k_z} is a function with the same periodicity as the crystal, i.e. $u_{k_z}(z+p) = u_{k_z}(z)$. There are different functions ψ corresponding to each wavevector k_z, which is why we write k_z as an index to ψ and u.

To find the unkown solutions k_z and ψ, we start by writing eq. 5.1 in a different way:

$$\psi_{k_z}(p) = e^{-jk_z p} \cdot \psi_{k_z}(0) \tag{5.2}$$

Once again, we expand the electromagnetic field in terms of the eigenmodes in each individual layer. We retain N modes in this series expansion. Since the total field consists of both a forward and a backward propagating component, we can represent ψ in $z = 0$ by a $2N \times 1$ column vector:

$$\psi \leftrightarrow \begin{bmatrix} F \\ B \end{bmatrix} \tag{5.3}$$

Here, F and B are $N \times 1$ column vectors representing the forward and backward propagating fields respectively.

We know from chapter 2 that we can construct a transfer matrix of the basic period that relates the forward and backward propagating fields on both sides:

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} \cdot \begin{bmatrix} F(z=0) \\ B(z=0) \end{bmatrix} = \begin{bmatrix} F(z=p) \\ B(z=p) \end{bmatrix} \tag{5.4}$$

Combining this with eq. 5.2, we get

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} \cdot \begin{bmatrix} F \\ B \end{bmatrix} = e^{-jk_z p} \cdot \begin{bmatrix} F \\ B \end{bmatrix} \tag{5.5}$$

This means that the vectors describing the field profiles are the eigenvectors of the transfer matrix. Using the eigenvalues λ of this matrix, we can calculate the Bloch vectors k_z from $\lambda = e^{-jk_z p}$[McRae1968]

[Gralak2000]. Because of the multi-valued nature of the complex exponential function, the k_z values are only defined up to an integer multiple of $2\pi/p$. We can therefore choose to restrict the values of k_z to the range $\pm\pi/p$, which also happens to be the range corresponding to the first Brillouin zone of a 1D periodic structure.

Since the transfer matrix is a $2N \times 2N$ matrix, we will find $2N$ Bloch modes. It turns out that half of these have $\Re(k_z) > 0$, the other half is the mirror image with $\Re(k_z) < 0$, so that we find only N modes in the irreducible Brillouin zone.

If for a given frequency (or wavelength), we find only evanescent Bloch modes (for which $\Im(k_z) \neq 0$), this frequency is said to lie within a *band gap*. In a band gap, no propagative modes can exist and any light incident on the structure from the outside will be completely reflected, because it can find no modes to couple to.

5.1.2 Inverting T_{21} using SVD

Let us recall the explicit form of the transfer matrix from eq. 2.46:

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}_{12} - \mathbf{R}_{21} \cdot \mathbf{T}_{21}^{-1} \cdot \mathbf{R}_{12} & \mathbf{R}_{21} \cdot \mathbf{T}_{21}^{-1} \\ -\mathbf{T}_{21}^{-1} \cdot \mathbf{R}_{12} & \mathbf{T}_{21}^{-1} \end{bmatrix} \tag{5.6}$$

We already mentioned that $\mathbf{T}_{21}$ can become singular in the presence of higher order modes. So, in order to calculate band diagrams, we somehow need to take care of this problem. We will therefore now discuss a method to invert nearly singular matrices using singular value decomposition (SVD).

SVD is the numerical applications programmer's best friend: not only does it diagnose the presence of singularities in matrices, in many cases it will also provide a useful numerical answer when performing operations on these matrices. SVD is based on the fact that any square complex matrix $\mathbf{A}$ can be written as [Press1993]

$$\mathbf{A} = \mathbf{U} \cdot \Sigma \cdot \mathbf{V}^H \tag{5.7}$$

Here, $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices, i.e. their hermitian conjugate is also their inverse ($\mathbf{I}$ is the unit matrix):

$$\mathbf{U} \cdot \mathbf{U}^H = \mathbf{I} \tag{5.8}$$

$$\mathbf{V} \cdot \mathbf{V}^H = \mathbf{I} \tag{5.9}$$

The matrix Σ is a diagonal matrix containing non-negative real numbers:

$$\Sigma = diag(\sigma_i) \tag{5.10}$$

These numbers are called the singular values of **A** and are ordered in decreasing magnitude, i.e. $\sigma_1 > \sigma_2 > ... > \sigma_N$.

Using this decomposition, we can easily form the inverse of **A** as

$$\mathbf{A}^{-1} = \mathbf{V} \cdot diag\left(\frac{1}{\sigma_i}\right) \cdot \mathbf{U}^H \tag{5.11}$$

When **A** is mathematically singular, one or more of the σ_i's will be identically zero and the previous equation will fail. If **A** is not mathematically singular, but numerically very close to it due to round-off errors, one or more of the σ_i's will be very small, on the order of the machine precision. This means that SVD can diagnose how close to singularity a matrix is. In this context, one defines the *condition number* of **A** as the ratio of the largest singular value to the smallest one:

$$C = \frac{\max_i(\sigma_i)}{\min_i(\sigma_i)} \tag{5.12}$$

The smaller C, the better **A** is conditioned numerically.

Not only can SVD diagnose ill-conditioned matrices, it can also be applied to give a useful numerical approximation of $\mathbf{A}^{-1}$. To do this, we zero the elements of $diag(1/\sigma_i)$ in eq. 5.11 when the magnitude of σ_i is smaller than some threshold, effectively replacing $1/0 = \infty$ by 0. This might seem like a strange approach at first, because this zeroing is equivalent to throwing away information present in the original matrix. However, it turns out that the information we throw away is precisely the one most affected by round-off errors. So, we keep only the 'stable' part of the problem and dispose of the 'unstable' part. Physically, this means throwing away some of the contributions of the higher-order modes, which are numerically unstable in the T-scheme because of their strong damping.

This engineer's explanation can of course be grounded more mathematically, but for that we refer the reader to [Press1993]. The gist of the matter is that eq. 5.11 constructs a kind of pseudo-inverse in the least-squares sense, because the columns of **U** span the range of **A**, and those of **V** its nullspace.

5.1.3 Example: DBR

As a first trivial example, we calculate the band structure of a DBR stack, a sequence of planar layers with a refractive index alternating between a high and a low value. For these two refractive indices, we take $n_1 = 3.4$ and $n_2 = 1$. If we design the structure for high reflectivity at $\lambda = 1$ μm, the layers should be a quarter wavelength long at $\lambda = 1$ μm, meaning lengths of $d_1 = 1/4n_1$ and $d_2 = 1/4n_2$. The period of the stack is therefore $p = d_1 + d_2 = 0.324$ μm in this case.

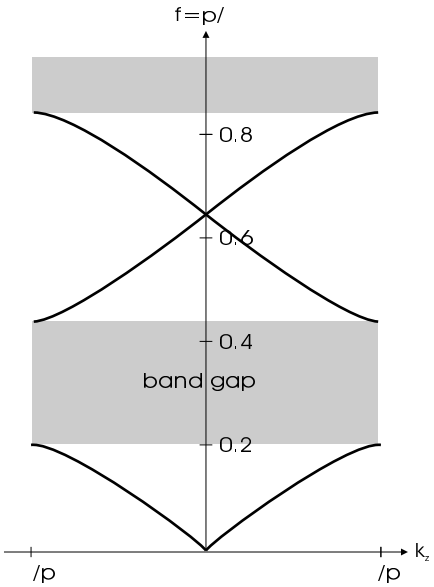

Figure 5.2: Band diagram of a DBR stack for normal incidence.

Fig. 5.2 shows the calculated band diagram of the DBR for normal incidence. The design wavelength of $\lambda = 1$ μm corresponds to a normalised frequency $f = p/\lambda = 0.324$. We clearly see that there is a band gap in this region. Light incident with a frequency in the gap will find no propagating modes to couple to, and will therefore be fully reflected.

5.1.4 Example: 2D square lattice

We will now try to use this method to calculate the band structure of a 2D square lattice of square dielectric rods with $n = 3.4$ in an air background (fig. 5.3). The ratio of the side of the square rods to the crystal period is $2r/a = 0.25$. This crystal is known to have a band gap for TE[1] polarised modes, i.e. when $\mathbf{E}$ is oriented parallel to the rods, perpendicular to the plane of the figure.

We will first concentrate on Bloch modes propagating along the ΓX crystallographic direction, indicated by z in fig. 5.3. We take a unit cell with length a as drawn on that same figure. For lateral boundary conditions, we place PEC walls parallel to the z-axis. These will act as mirrors and model the infinite extent of the crystal in the lateral direction. In order to get all modes, we also need to perform a seperate calculation with two PMC walls. The results from these two calculations are brought together in fig. 5.4. The k_z-axis is labeled with the

[1] In other works, e.g. [Joannopoulos1995], this is called the TM polarisation.

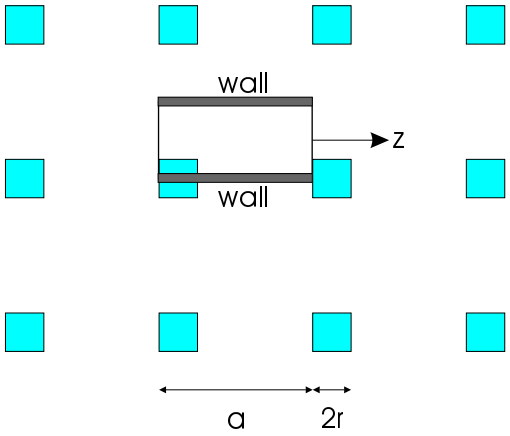

Figure 5.3: Square lattice of dielectric rods

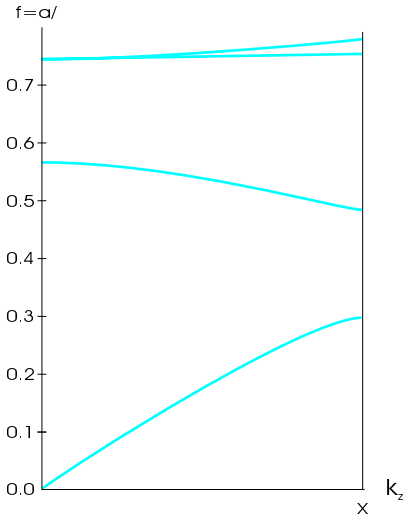

Figure 5.4: Band structure along ΓX for the TE modes of a square lattice.

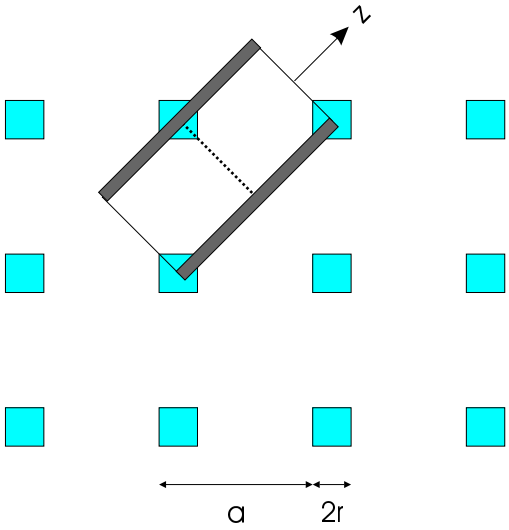

Figure 5.5: Square lattice with unit cell in ΓM direction

high-symmetry points of the Brillouin zone of the square lattice: $k_z = 0$ corresponds to the Γ-point, $k_z = \pi/a$ is the X-point. When we compare fig. 5.4 to e.g. the results in [Joannopoulos1995], we obtain a very good agreement. The band structures from [Joannopoulos1995] are calculated with the plane-wave method [Meade1993], which is based on variational principles and expands the fields in plane waves rather than in eigenmodes. Compared to the plane-wave method, the technique presented here seems to be very computationally efficient: the results from fig. 5.4 were obtained using a mere five eigenmodes, while plane-wave methods typically require hundreds if not thousands of plane waves.

Let us now consider modes that propagate along ΓM, i.e. under an angle of 45 degrees with ΓX. Fig. 5.5 shows the unit cell we have to construct in this case. To calculate the band structure along this direction, we can exploit the mirror symmetry of the unit cell along the dashed line. Both of these half unit cells have of course the same scattering matrices, but with the medium indices 1 and 2 interchanged. We made a staircase approximation of the rods, using 32 slices for half a unit cell.

From fig. 5.5 we also see that the unit cell has a period p of $\sqrt{2}a$. We know from the description in section 5.1.1 that our algorithm will yield propagation constants in the range $\pm\pi/\sqrt{2}a$. But let us now recall the layout of Brillouin zone of a square lattice in fig. 5.6. From this figure, we can deduce that the distance ΓM is $\sqrt{2}\pi/a$, which is twice as large as our maximum attainable k-vector $M_F = \pi/\sqrt{2}a$. Bloch theory then dictates that modes beyond M_F will be folded back in the range

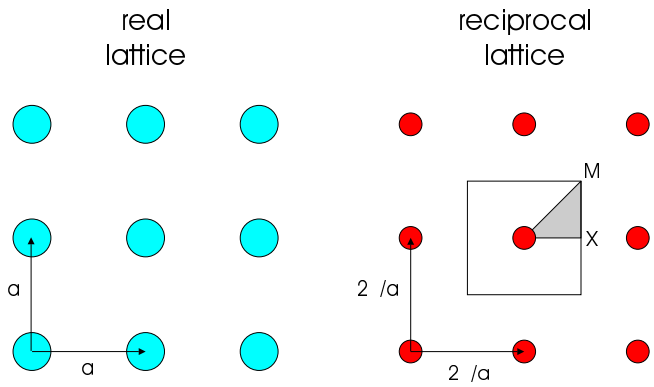

Figure 5.6: Reciprocal lattice and Brillouin zone of a square lattice.

ΓM_F (F as in 'Folded'), as indicated in fig. 5.7. For clarity, we have also plotted the unfolded band diagram, obtained simply by mirroring the bands between ΓM_F along the M_F-axis. As a guide to the eye, the 'normal' band structure (as would be obtained by e.g. the plane-wave method) is plotted in full lines, its mirrored copy in dashed lines. Our algorithm finds both the dashed and the full lines in the ΓM_F range. The fundamental reason for this complication is that there is not necessarily a connection between distances in a 2D reciprocal lattice and the reciprocal lattice of a 1D periodic structure we impose on a 2D periodic structure. This issue will be even more important when we study triangular lattices.

To end this section, we want to make a final remark with regard to the lateral mirroring walls that express the infinite nature of the crystal. One could rightfully ask why we choose the same type of wall on both sides of the cell. In fact, if we place a PEC wall on one side and a PMC wall on the other side, we do find modes, but they are of a different nature than the ones found with the same type of wall on both sides. We illustrate this with the transverse (x-direction) mode profiles sketched in fig. 5.8. For a TE mode, a PMC wall imposes a maximum of the electric field, while a PEC wall forces $\mathbf{E}$ to zero. In the case of PMCs on both sides of the unit cell, all cells have equal electric fields, both in amplitude and in phase. For the mixed PMC/PEC case, half of the cells have positive fields, the other half negative fields. This however results in modes with a lower symmetry: if we pick a row of cells at random, it is impossible to predict which of the two groups the row will belong to. For that reason, we do not include them in our band plots. They also do not show up in plane-wave calculations.

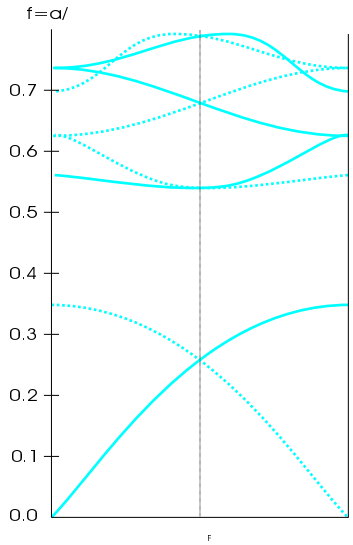

Figure 5.7: Band structure along ΓM for the TE modes of a square lattice.

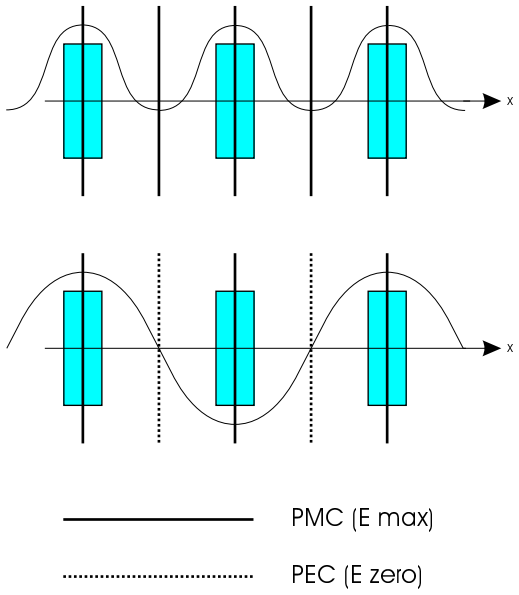

Figure 5.8: Influence of walls on symmetry of TE modes.

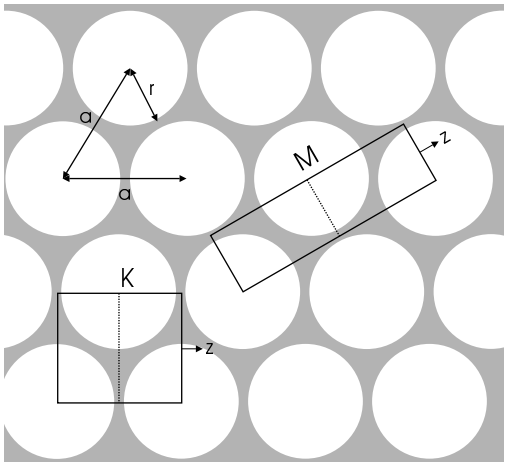

Figure 5.9: Triangular lattice of air holes.

5.1.5 Example: 2D triangular lattice

In this section, we will calculate the 2D band structure of the crystal from fig. 5.9. It is a triangular lattice of round air holes in a dielectric matrix with $\varepsilon_r = 13$. The ratio of the rod radius to the period is $r/a = 0.48$. This structure is known to possess a complete band gap, both for TE and TM polarisations and for any in-plane propagation direction.

Fig. 5.9 also shows the unit cells used to calculate the band structures in the ΓK and ΓM directions. The reciprocal lattice is drawn in fig. 5.10. It is a triangular lattice as well, but with period $2\pi/a$ and rotated over 90 degrees. From these figures, we can calculate that the distance ΓK is $2\pi/\sqrt{3}a$, while the maximum Bloch vector our algorithm will yield in that direction is only π/a. For the ΓM direction, these numbers are π/a and $\pi/\sqrt{3}a$ respectively. This means that the folding-over at K_F and M_F will occur at roughly 60% of the Brillouin zone, as shown in the actual calculated band structures (fig. 5.11).

For comparison purposes, we give in fig. 5.12 the band structure as computed by the plane-wave method [Joannopoulos1995]. It scans the full outer edge of the Brillouin zone, including the KM edge. Because K_F and M_F lie at about 60% of the edge of Brillouin zone, the folding-over process will not reach back to the Γ-point, as was the case in the previous section, but only to the points K′ and M′, indicated for the lowest order TM branches on fig. 5.11. However, between K′ and M′, the algorithm still finds branches which cannot be accounted for by the folding-over process. Careful comparison with fig. 5.12 however reveals that these branches behave just like the modes along KM, suggesting that they are formed by the projection of these modes along either the ΓK or the ΓM direction.

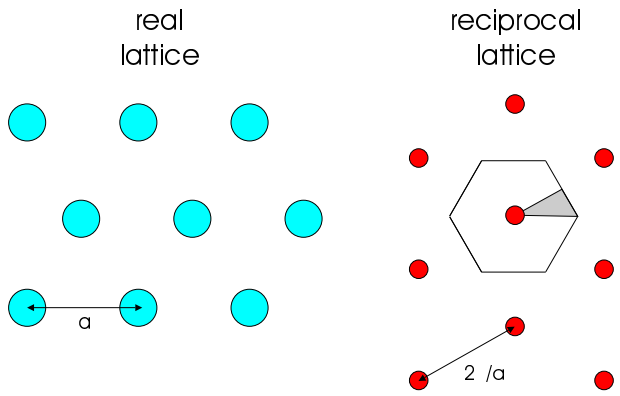

Figure 5.10: Reciprocal lattice and Brillouin zone of a triangular lattice.

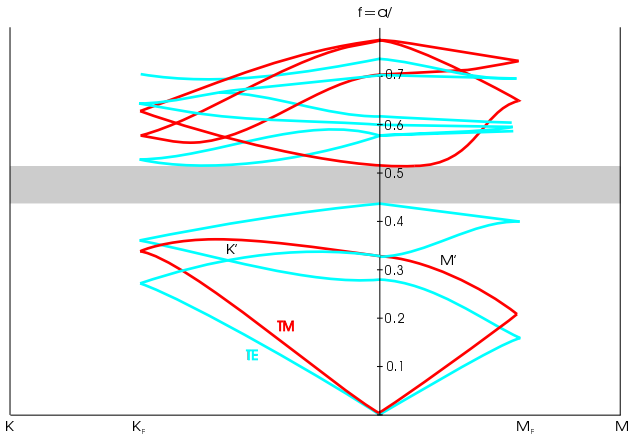

Figure 5.11: Band structure of a triangular lattice.

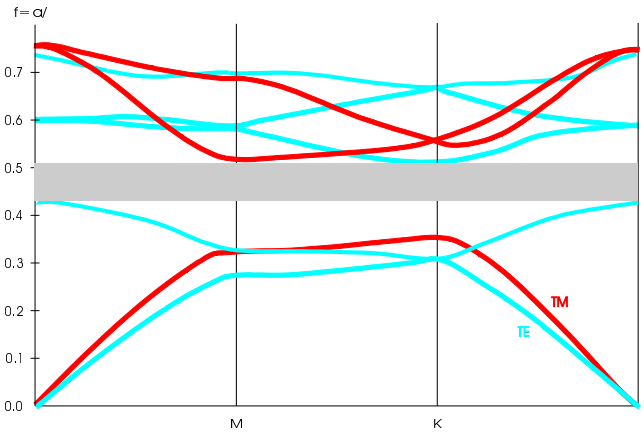

Figure 5.12: Band structure of a triangular lattice, as per [Joannopoulos1995].

So, in summary, the algorithm finds two kinds of branches for the triangular lattice case. One set of branches corresponds to the modes in the range ΓK_F or ΓM_F in the regular band diagram. The other set of branches has a different interpretation in the range $\Gamma K'$ (or $\Gamma M'$) than in the range $K'K_F$ (or $M'M_F$). In the former case, they correspond to the modes along KM in the regular band diagram. In the latter case, they are the folded-over modes that lie in the range $K_F K$ (or $M_F M$) in the regular diagram. Fig. 5.11 therefore contains exactly the same information as fig. 5.12, but presented in a different way.

Finally, we want to mention that the method presented here fails when we want to calculate the mode structure at points *inside* the Brillouin zone, for directions for which no suitable unit cell can be found. However, since we are mostly interested in the location of the band gap, which is adequately described by the *edges* of the Brillouin zone, this is no serious drawback.

5.2 Photonic crystal waveguides

If we introduce a line defect in a photonic crystal, e.g. by removing a row, we can create a photonic crystal waveguide (fig. 5.13). Light captured in this line defect finds itself surrounded on both sides by a highly reflecting crystal which it cannot penetrate, so it has no other option than to follow the line defect.

Fig. 5.14 shows the dispersion relation of the fundamental TE waveguide mode when we remove a row of rods from the square lattice of dielectric rods in air we studied in section 5.1.4. We see that this waveguide mode has a frequency inside the band gap, which is precisely

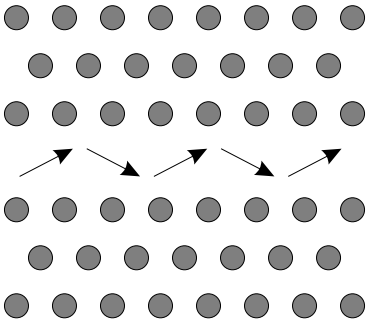

Figure 5.13: Photonic crystal waveguide.

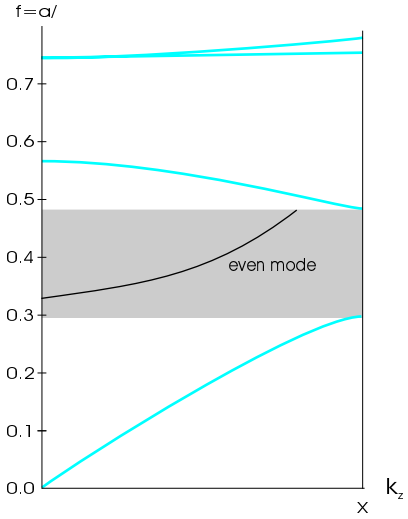

Figure 5.14: Band diagram of a photonic crystal waveguide, made by removing a row of dielectric rods in the ΓX direction in a square lattice.

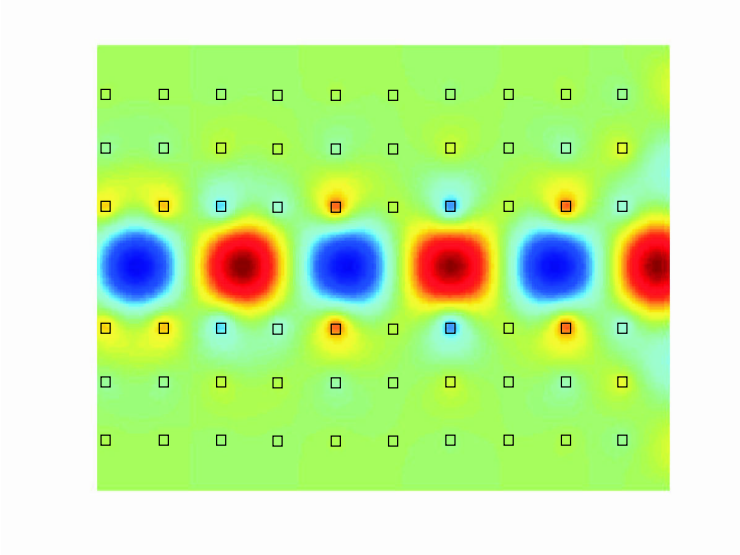

Figure 5.15: Field profile of a photonic crystal waveguide mode propagating in a finite crystal.

what one would expect because outside of the band gap the crystal is not highly reflecting. The structure is taken to be finite in the transverse direction, with three rows of rods on either side of the defect. Because this structure now has a 1D periodicity, the technique from section 5.1.1 to calculate the Bloch modes is useable without any complications.

In fig. 5.15, we show the electric field of the fundamental defect mode at $f = 0.4$ as it is propagating in a finite structure with 10 periods in the horizontal z-direction. We see that the mode is indeed confined to the line defect. It is quite remarkable that this waveguide guides light in air, rather than in a high index material as is the case in conventional waveguides.

It is worth pointing out that fig. 5.15 shows the *entire* computational domain used in the calculations. Other methods (e.g. the FDTD method) place a point source inside the waveguide to excite it. This requires a rather long stretch of waveguide to achieve an equilibrium field distribution containing only the desired mode. In eigenmode expansion methods things are different. Using the techniques from section 5.1.1, we can first calculate the expansion coefficients of the forward and backward fields that make up the Bloch mode in a waveguide that is infinite in the z-direction. We can then use the forward propagating component as the incident field to excite the left hand side of the finite structure in fig. 5.15. As we can see, this gives us a near equilibrium

field distribution immediately, without the need for a larger computational domain.

The explicit knowledge of the Bloch mode field profiles in the infinite structure also allows us to calculate quite easily the reflection coming from the right end facet of the waveguide, where it opens up into free space. FDTD calculates this reflection by sending a short pulse through the structure and by looking at the time evolution of the fields. In order to spatially distinguish the reflected pulse from the incident pulse, the computational domain has to be taken much larger than the actual structure of interest, which dramatically increases calculation times [Mekis1996]. Using eigenmode methods however, we can simply calculate the backward propagating field at the entrance of the finite structure and compare it to the backward field in the infinite structure. The difference between these two fields is then due to the reflections occurring in the finite structure. Because of standing wave effects, we have to average this reflection over several structures containing different numbers of periods in the z-direction. The end results of these calculations show that the termination of the waveguide reflects no more than 2% of the incident light. This is not surprising, given the fact that most of the light is trapped in an air core and does not have to cross a refractive index step in order to reach the semi-infinite air exit medium. This contrasts sharply with the 30% reflectivity at the end facet of a traditional semiconductor waveguide.

To illustrate the ability of CAMFR to deal with more complicated photonic crystal structures, we show in fig. 5.16 the field profiles in a 3dB splitter followed by a 90 degrees waveguide bend. In a very small area, the incident light is split into two equal fractions and guided around a very sharp bend. All this clearly shows the potential of photonic crystals when it comes to integrating optical functionality in an extremely compact way. The bends can be made within a very small area (a few square microns), which compares very favourably to traditional waveguide bends in PICs, which have a radius of curvature of about 1 mm.

We now give some execution times for this splitter simulation in order to quantify the speed of eigenmode expansion methods. These times were measured on a 250 MHz Sun UltraSparc processor, which corresponds to a medium-end PC by today's standards. Still, using 50 modes, it takes no more than two seconds to calculate the Bloch vectors and their field expansion coefficients of the infinite waveguide. Plotting the spatial field profile with a resolution of $a/20$ using the slow but accurate S-type excitation formalism takes about two minutes. Memory requirements are only a few megabytes. Similar calculations in FDTD would take at least an order of magnitude more resources, both in terms of time and memory. Moreover, for many simulations, the spatial field profile is not required. For example, when calculating the waveguide end-facet reflectivity, the knowledge of the expansion coefficients

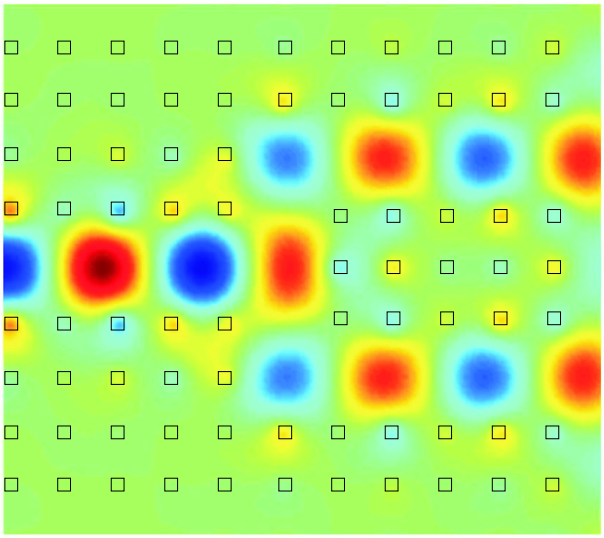

Figure 5.16: 3dB splitter in a photonic crystal.

is sufficient, while FDTD always needs to construct the full field profile and afterwards extract the reflection coefficient from it. In eigenmode methods, it is rather the other way around: a lot of useful information can be derived from the quickly calculated expansion coefficients, with the slighty slower spatial field profiles being optional.

5.3 Semi-infinite crystals

5.3.1 Reflection from semi-infinite crystals

Let us return to the eigenvectors of the transfer matrix we calculated in section 5.1.1. For simplicity, we will assume a 1D periodic medium consisting of layers that are uniform in the lateral x-direction. In this case, the reflection and transmission matrices of the basic period are diagonal, which means that in every layer the different modes (or different angles of incidence) are uncoupled. We can therefore consider each mode seperately, so that we can describe its reflection and transmission by scalars. This lightens our notational load, but the conclusions we will derive in this section are equally valid for the more general case involving matrices rather than scalars.

We already mentioned that the eigenvalues and eigenvectors of $\mathbf{T}$ can be divided into two groups according to the direction of propagation of the corresponding Bloch mode[2] (as determined by the sign of its

[2]In this section, it is important not to confuse Bloch modes with layer modes. Bloch

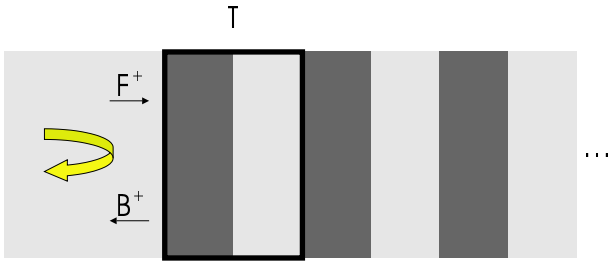

Figure 5.17: Reflection from semi-infinite crystal.

group velocity). In the case of uniform layers we consider here, $\mathbf{T}$ is a 2×2 matrix, with one eigenvector corresponding to a positive group velocity, the other to a negative one. We can collect the eigenvectors in a 2×2 matrix $\mathbf{E}$, where each column represents a separate eigenvector:

$$\mathbf{E} = \begin{bmatrix} F^+ & F^- \\ B^+ & B^- \end{bmatrix} \tag{5.13}$$

It is often found in literature [Nicorovici2001][Čtyroký1998] that the ratio B^+/F^+can be considered as the reflection coefficient of the forward layer mode at a semi-infinite repetition of the basic period (fig. 5.17). While this may be plausible at first sight, a thorough justification for this claim has to our knowledge never been presented. We will investigate this here in more detail, and at the same time show that this expression is not valid under all circumstances.

Suppose we have an arbitrary structure with the following transfer matrix:

$$\begin{bmatrix} F_2 \\ B_2 \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \cdot \begin{bmatrix} F_1 \\ B_1 \end{bmatrix} \tag{5.14}$$

The reflection of this structure as seen from medium 1 is $r = -C/D$, which can be seen by setting $F_1 = 1$, $B_1 = r$, $F_2 = t$ and $B_2 = 0$.

We will now calculate the reflection from a finite number n of repetitions of the basic period and try to take the limit as $n \to \infty$. We start by writing the transfer matrix $\mathbf{T}$ of the basic period in terms of its eigenvectors $\mathbf{E}$ and eigenvalues λ:

$$\mathbf{T} = \mathbf{E} \cdot \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \cdot \mathbf{E}^{-1} \tag{5.15}$$

Here, λ_1 correponds to the forward Bloch mode and can be written as $\lambda_1 = e^{-jk_z p}$. Similarly, $\lambda_2 = e^{+jk_z p}$.

modes are related to an infinite periodic structure, while layer modes are the eigenmodes of each $z-$invariant layer in the structure. In each layer, we expand the Bloch modes on the layer modes of that particular layer.

The transfer matrix $\mathbf{T}_n$ of n periods can be derived quite simply from eq. 5.15:

$$\mathbf{T}_n = \mathbf{T}^n = \mathbf{E} \cdot \begin{bmatrix} \lambda_1^n & 0 \\ 0 & \lambda_2^n \end{bmatrix} \cdot \mathbf{E}^{-1} \tag{5.16}$$

This is in fact another way of quickly calculating the scattering properties of a periodic stack (see section 2.4.3), but we prefer not to use it in practice because of the numerical instabilities associated with calculating $\mathbf{T}_{21}^{-1}$ for the determination of $\mathbf{T}$.

Using the explicit form of the eigenvectors from eq. 5.13, we can write

$$\mathbf{T}_n = \begin{bmatrix} F^+ & F^- \\ B^+ & B^- \end{bmatrix} \cdot \lambda_2^n \cdot \begin{bmatrix} (\lambda_1/\lambda_2)^n & 0 \\ 0 & 1 \end{bmatrix} \cdot \frac{1}{\det \mathbf{E}} \cdot \begin{bmatrix} B^- & -F^- \\ -B^+ & F^+ \end{bmatrix} \tag{5.17}$$

Since we are only interested in the ratio $-C/D$, we can factor out $\lambda_2^n / \det \mathbf{E}$ from eq. 5.17. If we assume that $\lambda_1 < \lambda_2$, which will be the case if the forward Bloch mode is damped exponentially, $(\lambda_1/\lambda_2)^n$ will tend to zero as $n \to \infty$. This means that

$$\mathbf{T}_\infty \propto \begin{bmatrix} F^+ & F^- \\ B^+ & B^- \end{bmatrix} \cdot \begin{bmatrix} 0 & 0 \\ -B^+ & F^+ \end{bmatrix} \tag{5.18}$$

In this case r_∞ is indeed

$$r_\infty = -\frac{C}{D} = -\frac{-B^- B^+}{B^- F^+} = \frac{B^+}{F^+} \tag{5.19}$$

However, for lossless Bloch modes k_z is purely real and $|\lambda_1/\lambda_2| = 1$. In this case, the limit $(\lambda_1/\lambda_2)^n$ as $n \to \infty$ is undetermined! This means that it is not possible to define the reflectivity from a semi-infinite stack under such circumstances. However, it turns out that $B^+/F^+$is still a good measure for the *average* reflectivity as a function of the number of periods n. This is illustrated in fig. 5.18, where we plot the amplitude of the reflection coefficient for the DBR from section 5.1.3 at a wavelength of 1.7 μm. The reflection exhibits a complicated, but periodic behaviour as a function of n. The average of this curve is well approximated by r_∞.

5.3.2 Semi-infinite crystals as boundary conditions

When simulating photonic crystal devices, we have so far only considered crystals with a finite number of periods in the lateral direction. However, it is tempting to try and devise a scheme whereby one could express that the crystal has in fact an infinite number of lateral periods. One approach of doing this would be to use boundary conditions to describe this semi-infinite crystal. This is illustrated in fig. 5.19, for

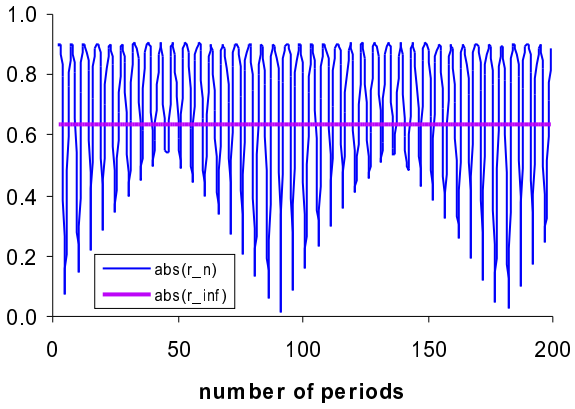

Figure 5.18: Reflection at a periodic structure as a function of the number of periods.

a single z-invariant waveguide that makes up an arbitrary 2D structure. For each angle of incidence, we can calculate the (average) reflection coefficient of such a crystal using the technique from the previous section. It is then straightforward to plug this information in the general slab waveguide dispersion relation from section 3.1. In this way, the semi-infinite crystal is reduced to a generalised boundary condition with an angle-dependent reflection coefficient.

This could have several advantages. First of all, the edge effects due to the finite lateral dimensions of the crystal could be eliminated. Secondly, the computational volume could be much smaller, which would translate in faster run times. Unfortunately, this technique has a number of drawbacks, which makes it unsuitable for practical purposes.

The first drawback is that the modes of this slab waveguide are not orthogonal anymore: we have seen in section 3.5.1 that the only angle-dependent boundary conditions that give rise to orthogonal modes are the transparent boundary conditions. Still, this is not an unsurmountable problem, because orthogonality is not required for the mode-matching procedure. Eq. 2.29 and 2.30 describe a $2N \times 2N$ linear system from which the $\mathbf{R}$ and $\mathbf{T}$ matrices can be determined jointly. Calculation times however will be longer than for the case of orthogonal modes, where we only have to solve an $N \times N$ system.

The second problem is altogether more serious, because we were unable to achieve convergence with this method, e.g. when we tried to find the Bloch modes of a defect waveguide in a triangular lattice. The most likely explanation for this is that the modes of the slab with these boundary conditions do not form a complete set. This is plausible from the interpretation of r_∞ we gave earlier. Under some circumstances, r_∞ is the *average* reflection of the crystal, because the oscillating nature

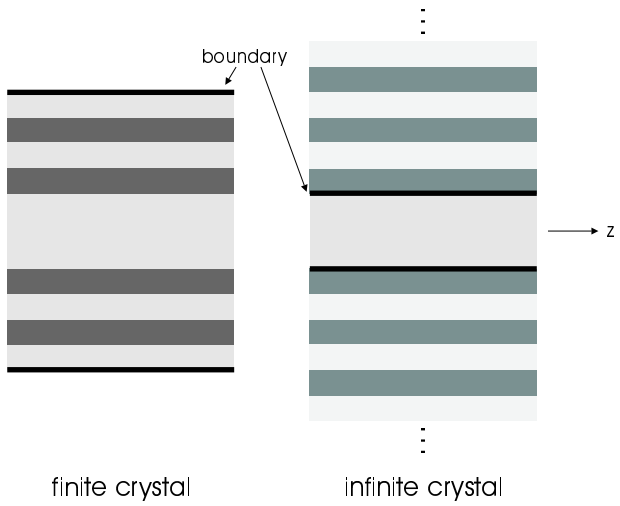

Figure 5.19: Boundary conditions describing semi-infinite periodic media.

of r_n makes it impossible to define a limit value for $n \to \infty$. It turns out that in these regions of the dispersion relation, no modes are found for the infinite crystal, while for finite crystals there are a number of modes in this region. For different values of n, these modes oscillate within this region without any limit value. This is illustrated in fig. 5.20, which compares the dispersion relations for an air core surrounded on both sides by the DBR from section 5.1.1. The wavelength is once again 1.7 μm and the air core is twice as thick as the air layer in the DBR. The dispersion relation for the semi-infinite crystal has no zeros in the region shown, while those for $n = 10$ and $n = 11$ periods clearly exhibit zeros in this range. These modes are obviously needed to describe the fields in finite crystals. The fact that they do not show up in infinite crystals is most likely what makes the set incomplete.

5.4 Other photonic crystal simulations

In this section, we briefly present some simulation results on photonic crystal related devices obtained by some of our colleagues using CAMFR. They form the core of their ongoing Ph.D. work, so we will only briefly discuss the main conclusions here.

Ir. Dirk Taillaert used CAMFR to design a surface coupling structure, which transforms a horizontally propagating waveguide mode into a vertical outcoupled field. This device has dimensions compatible with butt-coupling to a single mode fibre, and achieves a simulated coupling efficiency of over 90%.

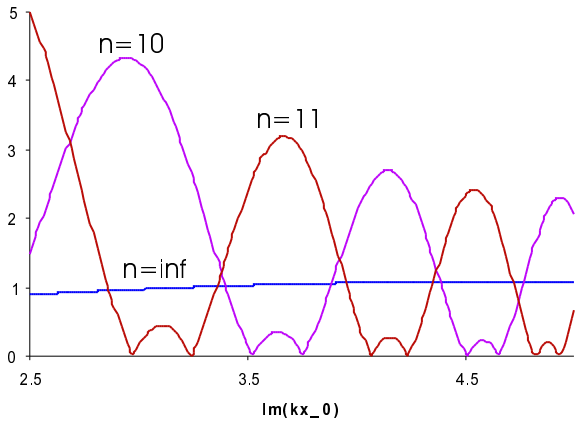

Figure 5.20: Dispersion relation for a core surrounded by DBRs with different numbers n of periods.

Ir. Wim Bogaerts studied the scattering losses that can occur in 1D photonic crystal slabs and investigated the influence of the refractive index contrast on these losses.

Ir. Bert Luyssaert analysed the properties of CROW waveguides, which are optical waveguides formed by coupled resonators. He is also evaluating novel techniques to use these waveguides to form compact tapers.

Finally, CAMFR also found its use to model waveguide bending losses using conformal transformations (Ir. Mathias Vanwolleghem, Erwin Six, Michael Roelens), to design a pickup head for a next-generation optical storage system (Ir. Frederik Fransoo) and to study ring resonators (ir. Ilse Christiaens). We will not discuss these activities here, because they have been started only recently, and because they are not related to photonic crystals.

5.4.1 Surface coupler

The idea behind the surface coupler is to design a photonic crystal structure that will allow light to butt-couple from an in-plane waveguide to a vertical single mode fibre (fig. 5.21). Compared to traditional edge couplers, such a surface coupler could have a number of advantages. First of all, it makes the PIC wafer-level testable, because no facets have to cleaved. Secondly, it could potentially have a larger coupling efficiency, if the photonic crystal is designed such that the outcoupled field is well-matched to the mode profile of the fibre.

As a first prototype, a 1D periodic structure was designed. It consists of a shallow second-order grating to perform the horizontal-vertical coupling (fig. 5.22). The grating was chosen to be shallow, in order to

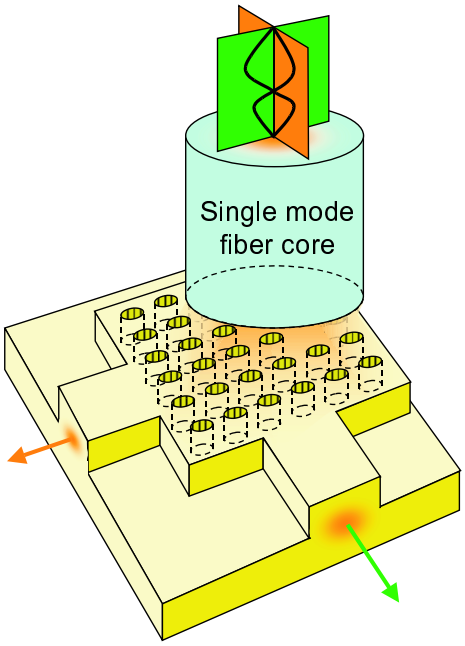

Figure 5.21: Surface coupler to single mode fiber.

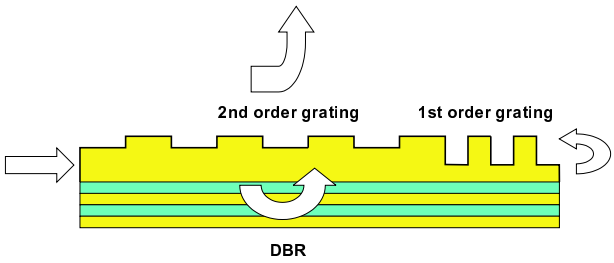

Figure 5.22: Sketch of 1D surface coupler.

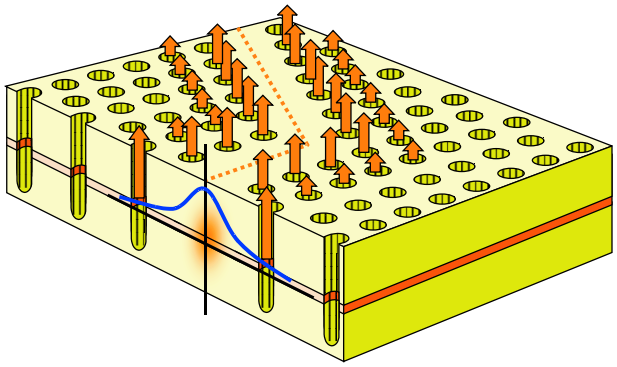

Figure 5.23: Photonic crystal slab.

reduce reflections when the light enters the coupler from the left. Light travels through the second-order grating section and is gradually coupled out. At the right, a strong first-order grating reflects the light and sends it back through the second-order grating for a second outcoupling pass. A DBR is placed underneath the entire structure, so that any light travelling down to substrate is reflected back up again.

The parameters of this structure were optimised using CAMFR, resulting in a compact 10 micron long device with a simulated outcoupling efficiency larger than 90%. We do want to point out that this is not a trivial structure to model, because light propagation in this device is completely non-paraxial. This means that it would be very hard to model this structure with the BPM method. Nor is classical eigenmode expansion with hard walls an option, because over 90% of the power is heading straight for the metal walls, where it would be completely reflected and disturb the simulation results. The FDTD method with PML boundary conditions is capable of handling such a structure, and in fact we did some verification runs against it, but its long calculation times make it rather unpractical for design work which requires a lot of iterations. We can therefore conclude that eigenmode expansion with advanced boundary conditions is a very powerful tool for these structures.

5.4.2 Scattering losses in photonic crystal slabs

3D photonic crystal structures are very difficult to fabricate. An attractive alternative therefore to achieve quasi-3D confinement of light is to combine a 2D photonic crystal with a slab waveguide (fig. 5.23). In these so-called photonic crystal slabs, light is confined in the vertical direction by the index contrast of the slab waveguide and in the in-plane direction by the photonic crystal, e.g. a lattice of etched air holes. However, when these air holes are etched through the slab, the

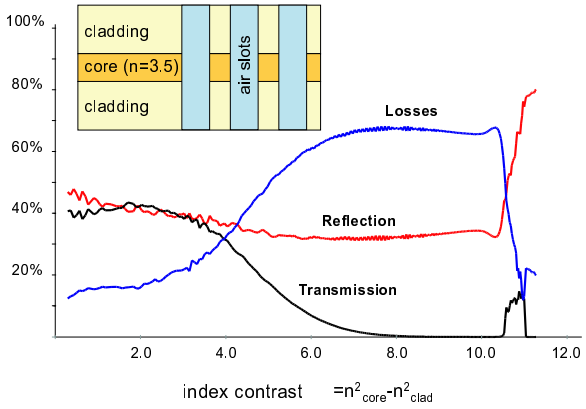

Figure 5.24: Radiation losses in a 1D photonic crystal slab.

vertical index confinement is essentially lost and light can be radiated out of the holes.

The study of these scattering losses is very important for practical applications. Specifically, it has been hotly debated whether one should go for high or low index contrast in these slabs to minimise the losses. One school of thought is based on an approximate theory to predict that these scattering losses will become more prominent as the index contrast increases [Benisty1999]. Others prefer working in high index contrast material systems [Chow2000], because then the possibility exists to excite lossless guided Bloch modes [Johnson1999].

Our rigorous simulations of 1D photonic crystal slabs show that in fact both points of view are valid. Fig. 5.24 plots the reflection, transmission and radiation losses of a 1D photonic crystal slab as a function of index contrast. For low index contrasts, we see indeed that radiation losses increase for higher index contrast, because the stronger index contrast entails more scattering. For high index contrasts however, we arrive in a regime where these losses drop dramatically, because of the existence of a lossless guided Bloch mode. This can be seen as a situation whereby all the scattering from the individual air holes interferes destructively, such that the mode itself is completely lossless.

All this is illustrated further in the band diagrams for these two regimes (fig. 5.25). An important feature in these diagrams is the so-called light line, which separates the guided Bloch modes from the radiation Bloch modes, which form a continuum in an open structure. As the index contrast increases, this light line shifts upwards. For low index contrast, we are working above the light line, so that light can couple easily to the radiation modes and scatter. For high index contrast, the light line shifts high enough to reveal a guided Bloch mode, which is totally lossless.

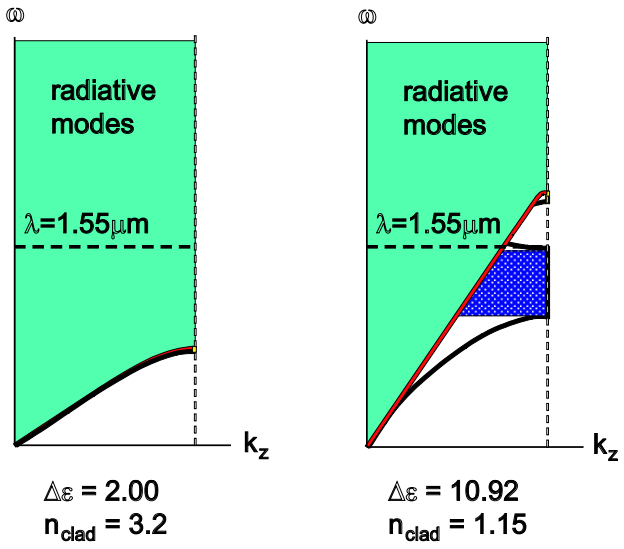

Figure 5.25: Band structure in photonic crystal slabs.

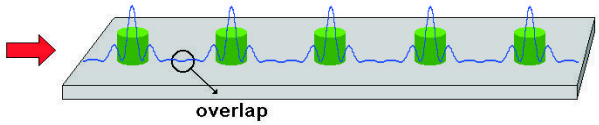

Figure 5.26: Coupled resonator optical waveguide.

5.4.3 CROW waveguides and tapers

Recently a new waveguide concept was introduced [Yariv1999], the Coupled Resonator Optical Waveguide or CROW. This waveguide consists of a series of resonators with a high Q-factor (fig. 5.26). Each of these resonators confines the light, but at the same time also lets a small part of it couple to the adjacent resonator. In this way, a lossless Bloch mode can be formed, not unlike the guided modes from section 5.4.2.

This type of waveguide has a number of advantages. Its dispersion relation is very flat, leading to a low group velocity. This can be an advantage in specific applications like second harmonic generation, because it slows down the light and increases the interaction. Also, by exploiting symmetries in the resonators, reflectionless bends can be constructed, as well as intersections of CROWs without any cross-talk [Xu2000].

Another application of CROWs currently under study at our department is its use as a very compact taper structure (fig. 5.27). Such

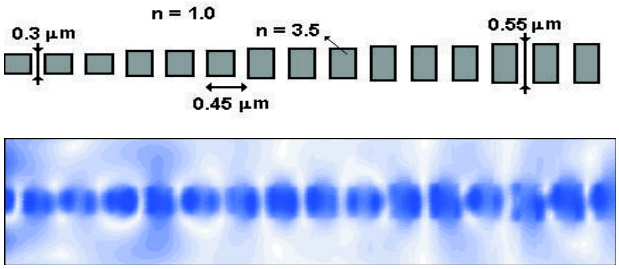

Figure 5.27: CROW taper.

a taper could be constructed by gradually changing the lateral dimension of the resonator, such that its mode broadens. Fig. 5.27 shows the field profile in a first prototype taper structure. This design is currently being optimised further.

Chapter 6

Vertical Cavity Lasers

> "Only two things are infinite: the universe and human stupidity, and I'm not sure about the former."
> Albert Einstein

This chapter is devoted to the modelling of laser cavities, in particular vertical-cavity surface-emitting lasers (VCSELs). We start by briefly reviewing the properties and advantages of VCSELs in section 6.1. In section 6.2, we will describe how we can extend CAMFR to find laser modes of an arbitrary cavity. We will use these techniques in section 6.3 to perform a detailed study on transverse optical confinement in VCSELs, which is a very important topic to achieve low-threshold high-performance devices.

6.1 VCSELs

VCSELs are lasers where the light propagates perpendicular to the wafer surface, rather than parallel to it, as is the case in traditional edge-emitting lasers. The cavity is formed by placing an active layer (like a quantum well or quantum dots) between two Distributed Bragg Reflectors (DBRs). These DBRs have to be highly reflecting, because the single-pass gain of these very thin active layers is rather small.

VCSELs have a number of advantages over edge-emitters. They can be made in cylindrical shapes, which results in a beam profile that is very well suited for coupling to optical fibre (fig. 6.1). Moreover, because they do not need cleaved facets as mirrors, they are wafer-level testable. On top of that, VCSELs can have very small cavity volumes, on the order of a few cubic wavelengths, which can result in very low threshold currents compared to edge-emitters [Huffaker1999]. Also because of their short cavity lengths, they inherently lase only in a single

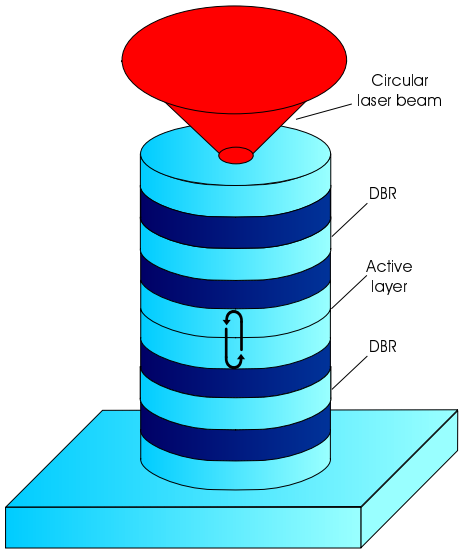

Figure 6.1: Vertical-cavity surface-emitting laser.

longitudinal mode, which opens up the promise of high speed modulation. One of their most promising aspects however is that they can be integrated into 2D arrays [Grabherr1999], because they emit light perpendicular to the substrate. This opens up the possibility of massively parallel optical data communication.

Contrary to photonic crystal devices, VCSELs are a commercial reality today, especially in the short infrared wavelength range (850 nm). At least a dozen companies worldwide fabricate these devices nowadays, and new start-ups are being announced on a regular basis. Commercial VCSELs are currently positioned in the market as low-cost alternatives to edge-emitting lasers in short to medium haul communications, where the requirements in terms of modulation bandwith and spectral purity are not as stringent as in long-haul communications. Especially the forthcoming introduction of Gigabit Ethernet systems might prove a real 'kill app' for VCSELs. Even now, there are already reports that the volume of VCSEL devices fabricated starts to outgrow that of edge-emitters.

The reason why today's best performing VCSELs emit in the short infrared, is that they can be fabricated in the AlGaAs-material system, where the growth technology is quite mature. Moreover, because of the high index contrast between GaAs and AlAs layers, DBRs with high reflectivity can be achieved with a relatively modest number of periods.

Things are rather less straightforward in other wavelength ranges, where VCSEL fabrication is still largely a research topic rather than a

commercial activity.

On the short-wavelength side, red VCSELs could be very attractive because that wavelength is very well matched to an attenuation minimum of plastic optical fibre (POF). These devices are fabricated in AlGaInP. Although this material system poses some problems with respect to thermal properties and carrier leakage, current results are promising [Onischenko2000] [Saarinen2000].

Another important wavelength range is long infrared, in particular the telecommunications wavelengths of 1.3 and 1.5 μm, which are until now mostly dominated by edge-emitters. Active layers at these wavelengths are fabricated in InP, but this material system can only provide low index contrast, which makes it difficult to fabricate high-reflectivity DBRs. Some approaches therefore combine an InP active layer with AlGaAs DBRs through a process called wafer-fusion [Ohiso1998]. Another method is to use quantum dots as the active material in the AlGaAs system. These dots can emit at longer wavelengths than quantum wells. Although only edge-emitters have been fabricated until now using this technology, the prospects for VCSELs are quite good [Maximov2000]. Finally, yet another material could be useful for long-wavelength emission in combination with the AlGaAs system, and that is the novel GaInNAs material [Larson1998].

There exist quite a number of VCSEL fabrication procedures. They differ mainly in how they realise transverse optical confinement, i.e. how they keep the laser beam together in the direction parallel to the substrate. Rather than describing these different VCSEL types here, we will introduce them more or less in chronological order in section 6.3, where we will compare their performances with simulation results. Another discussion of different VCSEL types can be found in the overview article [Giboney1998]. A more thorough treatment on all the aspects related to VCSEL design and fabrication is given in the textbook [Sale1995].

6.2 Finding laser modes

In this section we will outline how we can easily extend CAMFR to locate laser modes of arbitrary cavities. Once again, because of the abstract nature of the framework, these techniques apply equally well to cartesian as to cylindrical geometries. However, because most VCSELs have cylindrical symmetry (see e.g. fig. 6.1), all the examples in this chapter will be for cylindrical cavities.

A laser mode is a resonant optical field distribution that can propagate indefinitely back and forth inside the cavity without the need for any external sources. To formulate this mathematically, we start by dividing the cavity at an arbitrary location in a top and a bottom part (fig. 6.2). Using eigenmode expansion, we then calculate the reflection

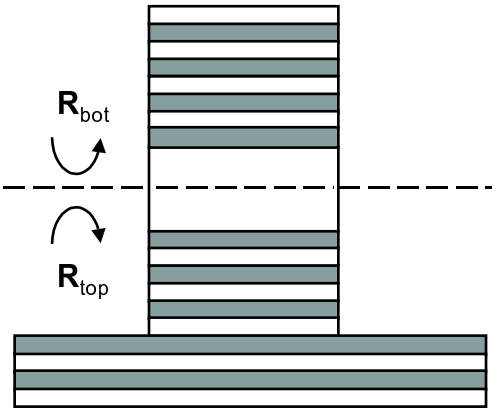

Figure 6.2: Generic laser cavity consisting of a top and a bottom part.

$\mathbf{R}_{top}$ of the top part as seen from the bottom, and similarly the reflection $\mathbf{R}_{bot}$ of the bottom part as seen from the top. Suppose we describe the laser mode by its (so far unknown) eigenmode expansion coefficients collected in the vector $\mathbf{A}_{lasing}$. A round trip of the field in the cavity consists of first reflecting off the bottom part and subsequently off the top part of the cavity. If this field profile is to propagate indefinitely in the cavity without the need for external sources, it should exactly reproduce itself after one round trip. This condition of unity round trip gain can be expressed as

$$\mathbf{R}_{top} \cdot \mathbf{R}_{bot} \cdot \mathbf{A}_{lasing} = \mathbf{A}_{lasing} \tag{6.1}$$

Put differently, if the cavity matrix $\mathbf{Q} = \mathbf{R}_{top} \cdot \mathbf{R}_{bot}$ has an eigenvector with an eigenvalue ν of 1, this eigenvector describes a lasing mode.

For an arbitrary structure, $\mathbf{Q}$ will not have such an eigenvector, so we will need to vary some parameters in order to achieve this resonance. First of all, we will need to vary the wavelength λ to get phase resonance, which is when the largest eigenvalue ν lies on the positive real axis. In the absence of any gain in the active region, ν will be smaller than unity, because of the losses inside the cavity. This means that we will also have to tune the material gain g_{mat} in the active region, in order to achieve amplitude resonance where the gain exactly compensates the losses. We do this by simply changing the imaginary part of the refractive index of the active layer, which is by definition related to g_{mat} by

$$g_{mat} = 2\frac{2\pi}{\lambda}\Im\left(n_{active}\right) \tag{6.2}$$

This quantity is usually expressed in units of cm^{-1}, and can be

thought of as the power gain of a plane wave $e^{-j(2\pi/\lambda)nz}$ after propagation over 1 cm in a medium with refractive index $n = n_{active}$.

In summary, locating a laser mode consists of doing a search in the two-dimensional (λ, g_{mat}) space to find a point where $\mathbf{Q}$ has an eigenvector with an eigenvalue of 1. This gives the laser mode's resonance wavelength, threshold material gain and field profile. Also note that the exact location of the gain region(s) with respect to the cavity cut is completely arbitrary.

To improve the numerical stability of this scheme, we could also make use of Singular Value Decomposition [Demeulenaere1999]. To show this, we start from the characteristic equation determining the eigenvalues

$$\det(\mathbf{Q} - \nu\mathbf{I}) = 0 \tag{6.3}$$

From this we can see that the matrix $\mathbf{Q} - \mathbf{I}$ becomes singular if $\mathbf{Q}$ has an eigenvalue $\nu = 1$. So, instead of looking for points where $\mathbf{Q}$ has an eigenvalue of 1, we can look for points where $\mathbf{Q} - \mathbf{I}$ has a singular value of 0. In this way, we enjoy all the stability advantages provided by SVD.

Finally, we want to point out that in practice, we do not really need to perform a search in a 2D space to achieve resonance. It turns out that in the vast majority of cases, we can suffice by doing a sequence of 1D searches, as is illustrated in fig. 6.3. The VCSEL structure in this example is a thin-oxide VCSEL (see section 6.3.4). Keeping the gain fixed at zero, we first vary the wavelength until we arrive at a minimum in the smallest singular value. Subsequently, we keep the wavelength fixed at this minimum and increase the material gain until we locate the laser mode. If we need higher precision, we can continue to vary wavelength and gain around this point, but usually, we can get a singular value on the order of 10^{-4} with one wavelength and one gain sweep. Locating the minimum is done numerically with Brent's algorithm [Press1993], which makes use of quadratic extrapolation to reduce the number of iterations required. The dots in fig. 6.3 indicate the actual iterations performed. Clearly, more iterations are needed in the neighbourhood of the minimum, in order to locate it precisely. Using 100 modes, the entire process of finding the wavelength and gain of this laser mode only takes about five minutes.

6.3 Transverse confinement in VCSELs

To illustrate the model presented in the previous section, we will spend the remainder of this chapter studying the important topic of transverse confinement in VCSELs. Transverse confinement deals with keeping the laser beam from spreading out too much in the in-plane direc-

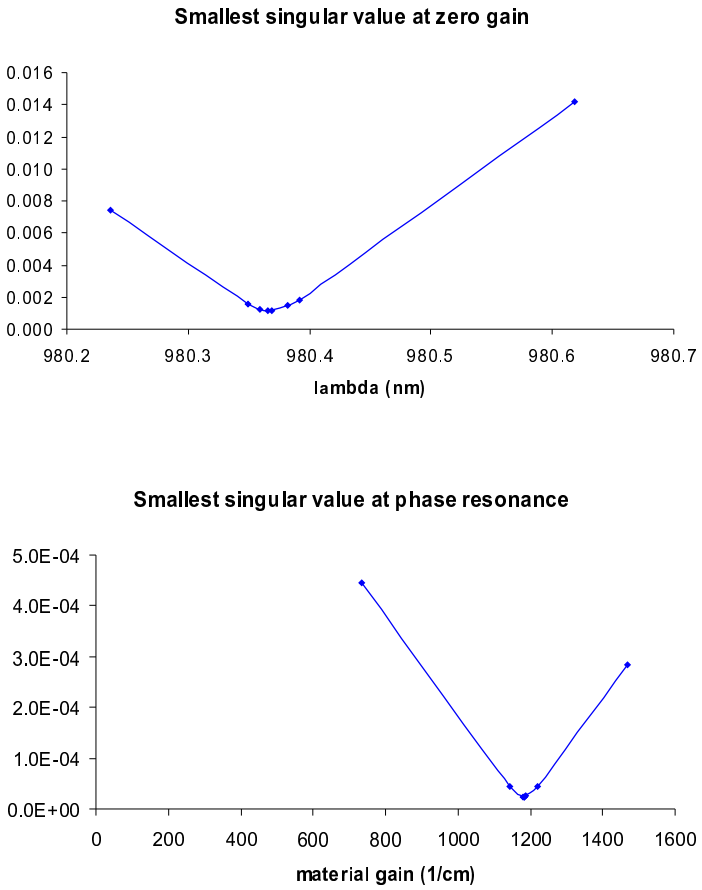

Figure 6.3: Varying wavelength and gain in order to achieve a zero singular value.

tion, so that it remains focused in a narrow area in the active region. In this way, we can create lasers with very small lateral dimensions. If we are also able to inject current only into this focus spot in the active region, these lasers will have very low threshold currents.

Over the past few years, several approaches to achieve transverse confinement have been tried out experimentally, some of which are now well-established, others are more speculative and research-oriented. We will compare these approaches numerically, which will show among others that designing transverse confinement in VCSELs is not a trivial matter, because it involves some trade-offs and because its design cannot be separated from that of the rest of the cavity.

We will also have a chance to compare our optical VCSEL model to a significant fraction of the other existing VCSEL models through a simulation comparison we organised [Bienstman2001a]. This will shed some light on the validity of approximations some other models make and on the necessity of the vectorial treatment of the optical fields.

6.3.1 Proton-implanted VCSELs

Proton implantation achieves confinement by selectively bombarding areas of the substrate with protons. This destroys the lattice structure of the semiconductor, turning the exposed regions into an isolator. By leaving circular apertures free from implantation, conducting regions surrounded by isolators can be created, which funnel the injected current into a small area in the active region [Lear1994].

Proton implantation does not change the optical properties of the exposed regions appreciably, so the only refractive index change in the VCSEL is the localised gain in the active region below the aperture due to current injection. For this reason, proton-implanted VCSELs are also called gain-guided devices, because only the gain region inside the active layer keeps the laser beam confined laterally. This is however a very small confining effect, so thresholds in these devices tend to be relatively high.

To illustrate this, we calculated the threshold material gain as a function of aperture size for the VCSEL layer structure from table 6.1. It is an AlGaAs VCSEL designed for emission around 980 nm. The cavity is one optical wavelength long, has 29.5 bottom DBR pairs and 25 top DBR pairs. It has a single 5 nm thick quantum well (QW) as an active region. The gain profile is taken to be piecewise constant. Inside the aperture, there is a constant gain, with a level to be determined by the laser mode locating algorithm. Outside the aperture, there is a small constant loss, indicated by an imaginary refractive index of -0.01. This structure is not chosen to mimic a realistic VCSEL design in all of its details, but rather as an illustration of the fundamental optical processes at work in VCSELs.

	Thickness (nm)	Material	Refr. index
air		air	1.00
25 pair DBR	69.49	GaAs	3.53
	79.63	AlGaAs	3.08
	136.49	GaAs	3.53
lambda cavity	5	QW	3.53+...j
	136.49	GaAs	3.53
29.5 pair DBR	79.63	AlGaAs	3.08
	69.49	GaAs	3.53
substrate		GaAs	3.53

Table 6.1: Layer structure of a proton-implanted VCSEL.

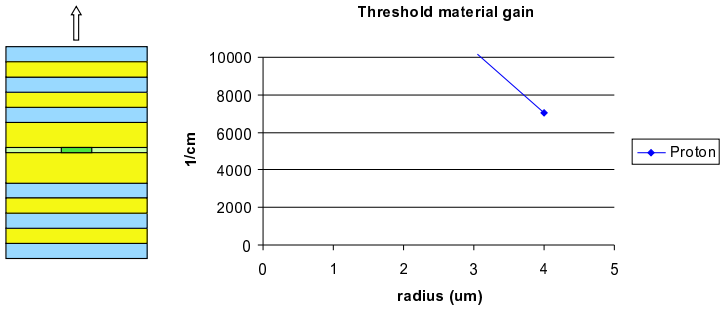

Figure 6.4: Threshold material gain for the proton-implanted VCSEL.

We can see from the results in fig. 6.4 that the threshold material gain quickly increases for decreasing aperture sizes. The results are plotted for the fundamental mode, i.e. the HE_{11} mode (the lowest order mode with Bessel order 1). The trends from this figure can easily be explained as follows. As the device size decreases, diffraction effects will become more prominent and will cause the laser beam to spread out. After reflecting at the DBRs, the mode will have a size much larger than the pumped gain region, and this poor overlap of the optical mode with the gain profile will lead to a very inefficient amplification process of the laser beam.

Because the gain-guiding offers only a very weak confinement, these *diffraction losses* will quickly lead to very high thresholds for small devices. As a rule of thumb, a single QW can provide around 500 cm^{-1} gain and typical VCSELs have 3 to 5 QWs. From fig. 6.4 we can therefore conclude that VCSELs with small diameters are not really feasible.

Until very recently, all commercially available VCSELs were of this type, and they still form a large part of the devices on offer. Typical characteristics of these devices are threshold currents of a few mA and optical output powers of a few mW.

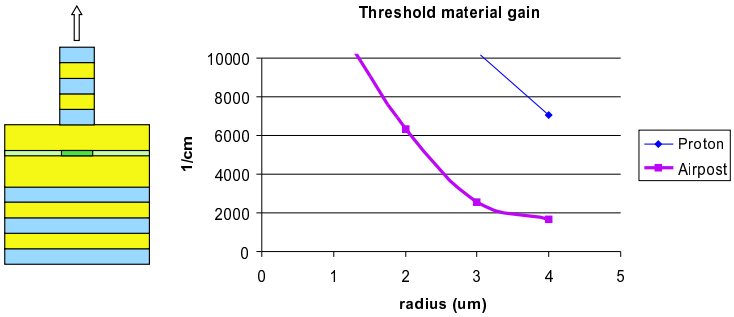

Figure 6.5: Threshold material gain for the airpost VCSEL.

6.3.2 Airpost VCSELs

A second approach to achieve confinement in VCSELs is etching away the semiconductor material around the cavity to leave an "airpost" standing (i.e. a post surrounded by air)[Geels1991]. Typically only the top DBR is etched away, because etching through the active layer will create centres for non-radiative recombination which will lower the efficiency.

The etched post is in effect a waveguide, which will concentrate the optical field because it contains high-index material. This provides a much stronger optical confinement than pure gain guiding. We therefore expect the diffraction losses to be lower.

This is illustrated by the simulation results from fig. 6.5. The 'background' VCSEL structure is the same as from table 6.1, to allow easy comparison between the different forms of optical confinement. The gain diameter is taken to be the same as the post diameter. Clearly, the airpost VCSEL has much lower thresholds, making lasing in devices with a radius down to 3 μm feasible.

6.3.3 Thick-oxide-confined VCSELs

A technique that is currently very popular for lateral confinement in VCSELs is the use of aluminium oxidation. This process was first described in [Dallesasse1990] and consists of etching an airpost to expose a semiconductor layer with high Al content, like e.g. $Al_xGa_{1-x}As$ with the Al fraction x larger than 90%. The sample is subsequently introduced in a furnace at a temperature of around 400 degrees, through which a steady flow of water vapour is maintained. This hydrogen-rich environment will laterally oxidise the exposed AlGaAs layer, i.e. starting from the perimeter of the airpost and then working its way in, turning the AlGaAs into AlOx.

If this process is controlled accurately and stopped at the right time,

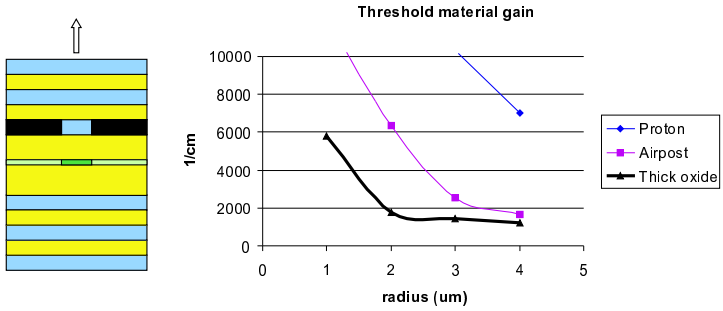

Figure 6.6: Threshold material gain for the thick oxide aperture VCSEL.

one can create a VCSEL where one layer consists of a central unoxidised conducting AlGaAs layer and an outer ring of isolating AlOx. This creates a current aperture, that can be used to restrict current injection to a small region in the active layer. Moreover, AlOx has a refractive index of around 1.55, which means that the oxide aperture will also act as a waveguide counteracting the diffraction losses.

Fig. 6.6 plots the threshold material gain of the reference VCSEL from table 6.1, but this time with the bottom layer of the top DBR replaced by a partially oxidised $\lambda/4$ thick AlGaAs layer. The gain diameter is taken to be the same as the oxide diameter. We make abstraction of the airpost that was etched to expose this layer, because this post typically has a diameter much larger than that of the oxide aperture itself. It will therefore not affect the optical field appreciably.

From fig. 6.6 it is apparent that this oxide aperture performs even better than the airpost VCSEL. At first sight this might seem surprising, because the oxide clearly has a much smaller confining effect than the airpost: the confining effect is only present in a single layer and the refractive index contrast is lower. So, there clearly is a mechanism other than diffraction loss at work in these devices.

To illustrate this, we compare the field profiles in an oxide-confined and an airpost VCSEL with a radius of 1 μm. These figures show the magnitude of the E_ϕ component of the optical field on a logarithmic scale. The left-hand edge of the figures is the symmetry axis $\rho = 0$, meaning that only the right half of the cavity is shown. Although the airpost achieves much better confinement in the top DBR, a lot of the light is lost at the corner interface between the unetched region, the etched airpost and the surrounding air layer. There, the light radiates out of the cavity in a wave travelling to the top right of the picture. These *scattering losses* are typical when light encounters a sharp refractive index step. For the oxide aperture, the weaker index contrast results in reduced scattering losses. At the same time however, the confinement is not as good, leading to higher diffraction losses. In this

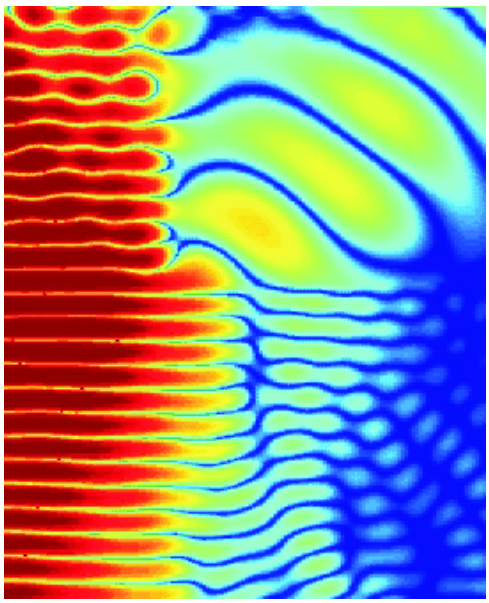

airpost VCSEL

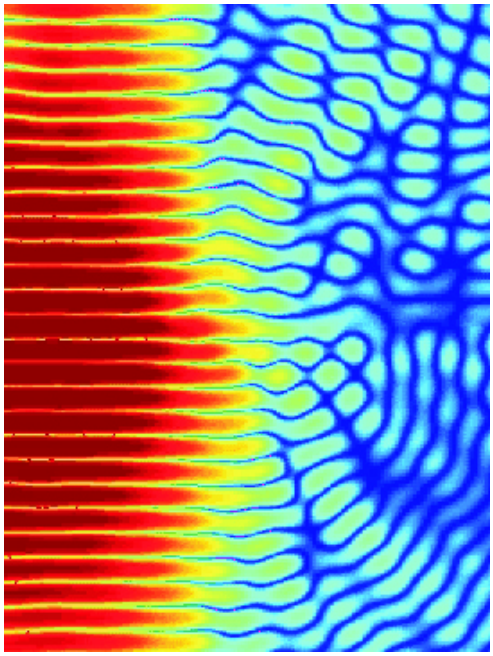

thick oxide VCSEL

Figure 6.7: Field profiles in an airpost VCSEL and a thick oxide VCSEL (the left edge of the figure is the rotational symmetry axis).

case, the balance between these two effects is such that the oxide-confined VCSEL has a lower threshold material gain than the airpost VCSEL.

In general, we can say that introducing any confining structure into a VCSEL will have two effects, one positive and one negative. On one hand, it will keep the optical field together when it is *inside* the confining structure, leading to reduced diffraction losses. On the other hand, coupling light into and out of this confining element disturbs the field and can lead to scattering losses at the *edges* of the structure. These effects tend to oppose each other in the sense that structures providing tight confinement also exhibit rather high scattering, so there is a fine balance to be struck.

6.3.4 Thin oxide-confined VCSELs

6.3.4.1 Aperture position dependent effects

This trade-off between diffraction and scattering losses leads to the reasoning behind the use of oxide apertures that are thinner than a quarter-wavelength layer. The fundamental question here is whether we can sacrifice some optical confinement by reducing the thickness of the oxide layer. This could be compensated by a reduction in scattering loss since the thinner oxide will disturb the optical field less.

By reducing the thickness of the oxide layer, we also gain an additional degree of freedom, namely the position of the aperture with respect to the optical field. Fig. 6.8 illustrates how the aperture can be placed at either a maximum (antinode) or a minimum (node) of the optical field.

In fig. 6.9, we show the threshold material gain for a node and an antinode aperture VCSEL with a thickness of $\lambda/20$, i.e one-fifth of the quarter-wavelength layer. Once again, the background VCSEL structure is the same. For the antinode oxide, the thresholds are the lowest of the devices we studied so far. This means that the reduction in scattering losses outweighs the small loss of confinement. For the node oxide, the situation is very different. Because it is placed at a field minimum, it will hardly have any effect on the optical mode, meaning low scattering but also very poor optical confinement. The latter effect is much stronger in this case, so the thresholds of this device are quite high.

To illustrate this further, we plot the field profiles in the antinode and the node device on the same scale (fig. 6.10). The poor confinement and the spreading of the mode is clearly visible in the node case.

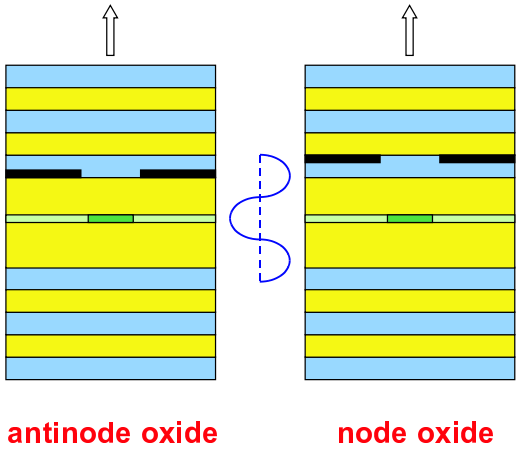

Figure 6.8: VCSEL with a thin oxide aperture at an antinode or a node position.

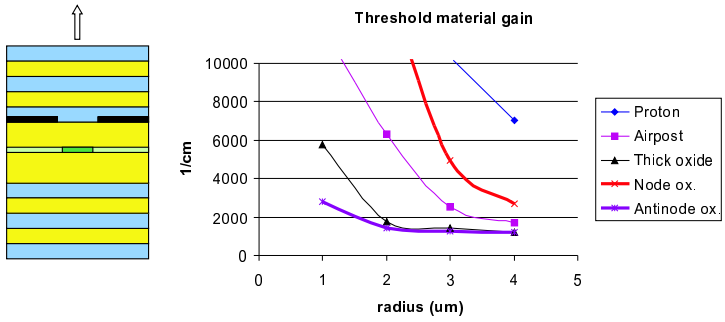

Figure 6.9: Threshold material gain for the thin oxide VCSEL.

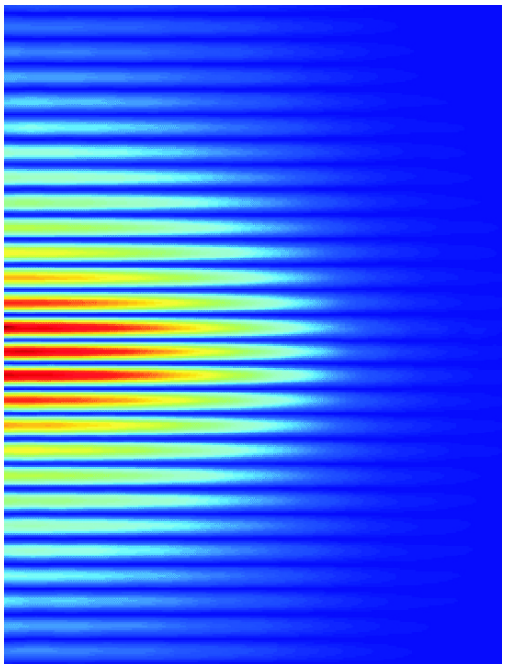

antinode oxide

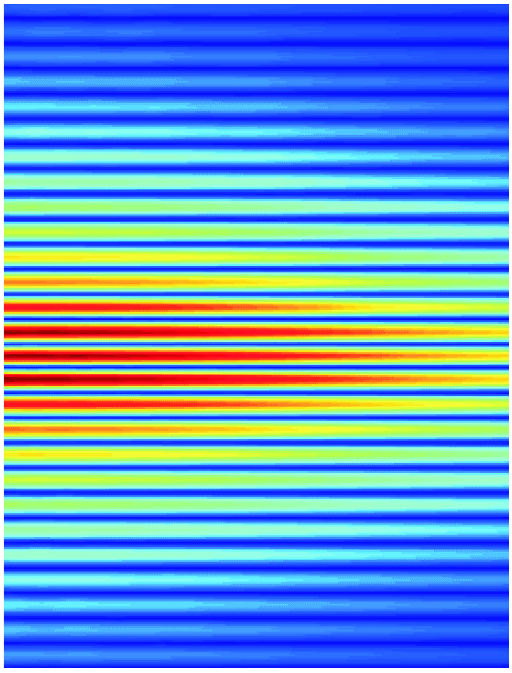

node oxide

Figure 6.10: Field profiles in VCSELs with antinode and node aperture ((the left edge of the figure is the rotational symmetry axis).

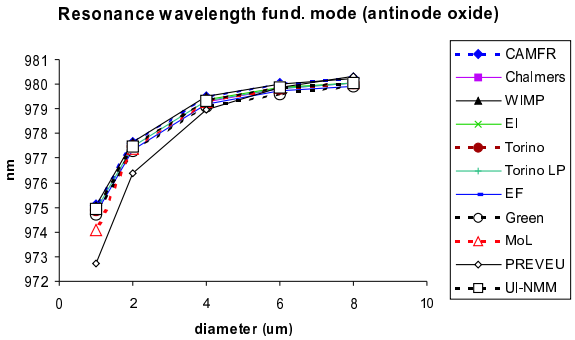

Figure 6.11: Resonance wavelength of the thin antinode oxide VCSEL.

6.3.4.2 Comparing different optical VCSEL models

We chose this topic of thin oxide apertures as a benchmark to compare different optical VCSEL models [Bienstman2001a]. We coordinated this activity in the framework of the European COST268 action, but many non-European research groups also participated.

Over the years, a number of approaches have been proposed to model VCSELs, either scalar or vectorial, or approximate or rigorous. However, comparing these models and therefore validating their underlying assumptions has been difficult if not impossible, since the published results obtained with these models are all for different VCSEL structures. It was the aim of the paper [Bienstman2001a] to compare many of today's important VCSEL models on the same benchmark problem, thereby quantifying the spread of results among the different models.

One of the benchmark problems was to calculate the properties of the thin antinode oxide VCSEL from section 6.3.4 as a function of the aperture size. From fig. 6.11, we see a decrease in lasing wavelength for decreasing aperture size. All by all, the difference between the models is mostly relatively small, no larger than a few percent. For a brief overview of the working principles of each of these models, we refer to [Bienstman2001a] and the references therein. Interesting to note however is that the scalar (approximate) models are represented by full lines in the figures, the vectorial (exact) models by dashed lines.

The differences between the models becomes larger when we plot the threshold material gain in fig. 6.12 (note the logarithmic scale).

Obviously, reducing the aperture size leads to a loss in confinement and higher thresholds. Differences between the models are larger here, about 10%. Also, there is a clear clustering between scalar and vectorial models, with the vectorial models consistently predicting higher thresholds. This is to be expected, since these models are able to take diffraction losses into account, which the scalar models cannot, be-

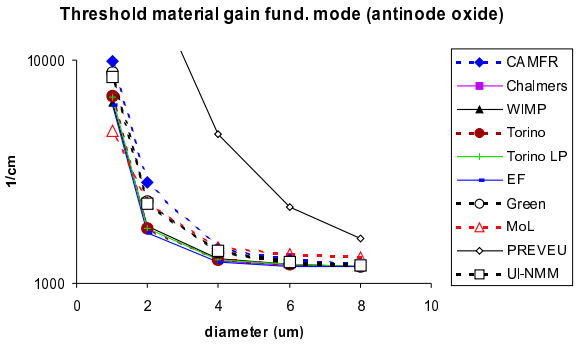

Figure 6.12: Threshold material gain of the thin antinode oxide VCSEL.

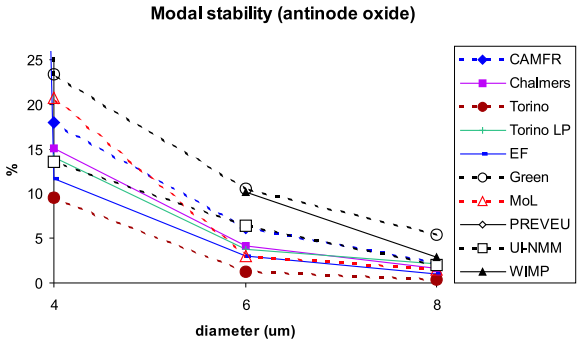

Figure 6.13: Modal stability of the thin antinode oxide VCSEL.

cause most of them are based on an effective-index approach. It is also clearly visible that the differences in predicted threshold become larger as the diameter is reduced.

An even larger difference between the models is apparent when we also consider the first-order mode (TE_{01} for the vectorial models, and LP_{11} for the scalar models). In fig. 6.13, we plot the modal stability, as defined by the percentual difference between the threshold g_0 of the fundamental mode and that of the first-order mode g_1:

$$\text{modal stability} = \frac{g_1 - g_0}{g_0} \tag{6.4}$$

As expected, going to smaller diameters increases the modal stability, because the first-order mode starts to approach cut-off.

In summary, we can say that the differences between the models gets larger as one moves away from paraxial propagation and into the regime of badly confined modes, as happens e.g. when decreasing the device diameter or modelling higher order modes. For a more in-depth discussion, we refer the reader to [Bienstman2001a].

	Thickness (nm)	Material	Refr. index
air		air	1.00
28 pair top DBR	69.8	GaAs	3.51
	81.4	$Al_{.87}Ga_{.13}As$	3.01
	349.0	GaAs	3.51
node oxide	30.0	AlOx/AlAs	1.55/2.93
	30.0	$Al_{.27}Ga_{.73}As$	3.36
QW	8.0	InGaAs	3.51+...j
	30.0	$Al_{.27}Ga_{.73}As$	3.36
	120.0	$Al_{.50}Ga_{.50}As$	3.22
etch stop layer	558.4	GaAs	3.51
31.5 pair bottom DBR	83.6	AlAs	2.93
	69.8	GaAs	3.51
substrate		GaAs	3.51

Table 6.2: Layer structure of the node oxide VCSEL from [Bond1999].

The comparison also allowed to get some feeling for the different run times of the vectorial models, like e.g. [Klein1998] [Noble1998], with CAMFR clearly performing better.

6.3.4.3 Influence of carrier diffusion

In the VCSEL structures we modelled so far, the antinode oxide always outperformed the node oxide in terms of threshold material gain. However, in some reports in literature, we can find devices where exactly the opposite behaviour is seen experimentally. One example of this is a study performed at the University of Southern California at Los Angeles (USC) [Bond1999]. In this section, we will try to elucidate these apparent differences between theory and experiment.

The USC device has a very different structure than the COST VCSEL, so it is not adviseable to extrapolate its behaviour from that of the COST device, at least not intuitively without rigorous simulations. The USC layer structure is given in table 6.2, and the main differences with the COST structure are illustrated in fig. 6.14.

The device has a cavity length of 4λ, rather than the short 1λ devices we considered so far. Diffraction losses will therefore play a larger role in these longer cavities. Additionally, the USC node oxide is placed at the first field minimum, whereas the COST node aperture is placed at the second field minimum (fig. 6.8).

Fig. 6.15 shows the simulated thresholds for the USC devices. The relative threshold difference between the node and the antinode devices is smaller than in fig. 6.9, but still, the antinode device has a lower threshold, contrary to the experimental evidence. The difference

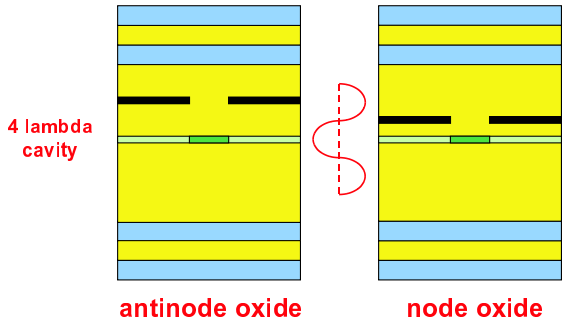

Figure 6.14: VCSEL structure from [Bond1999].

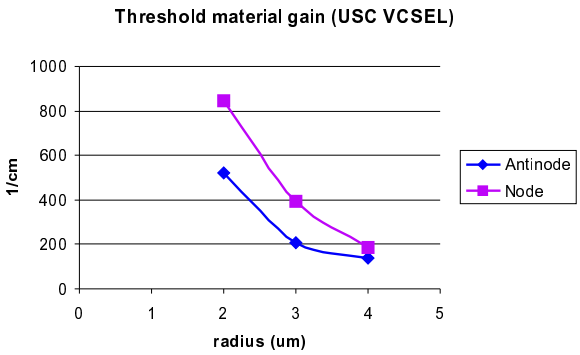

Figure 6.15: Threshold material gain the the USC thin oxide VCSELs.

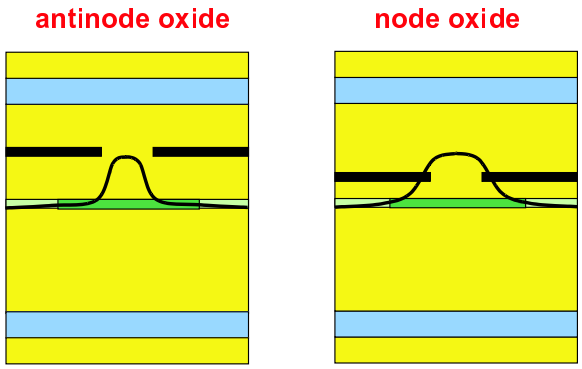

Figure 6.16: Influence of lateral carrier diffusion on the overlap of the gain profile with the optical mode.

in cavity design is therefore not sufficient to explain the observed experimental trends.

Clearly, another effect is important here, and that effect turns out to be carrier diffusion. Once carriers get injected through the oxide aperture inside the quantum wells, they tend to diffuse laterally. This means that the effective gain region will have a diameter larger than the oxide aperture diameter (fig. 6.16). We know from fig. 6.10 that the lower optical confinement in the node oxide will cause the laser mode to spread. This larger modal cross-section is very well matched to the larger gain diameter caused by carrier diffusion (fig. 6.16). For the antinode oxide, the opposite is true. The mode profile is rather narrow, meaning that a significant fraction of the gain in the QW only sees the tail regions of the optical mode. Because of the low field intensity in these tails, pumping is very inefficient there and this fraction of the gain is effectively wasted.

We can quite easily simulate the effects of carrier diffusion by choosing the gain diameter larger than the oxide aperture diameter. The precise size of the gain region should follow from other non-optical models (carrier diffusion models), so we will just treat the gain diameter as a phenomenological parameter here and simulate the device for a number of values of the gain diameter. Fig. 6.17 shows the results of such simulations for devices with a 2 μm radius oxide aperture. We can clearly see that the node oxide starts to outperform the antinode oxide from a gain radius of 3 μm, i.e. a diffusion length of about 1 μm. These values are in the same range as those reported in [Bond1999], which indicates that lateral carrier diffusion is indeed the relevant effect explaining the experimental trends.

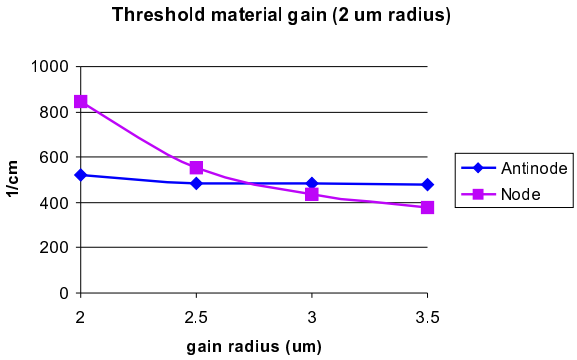

Figure 6.17: Threshold material gain for the USC VCSELs with a 2 μm oxide aperture radius.

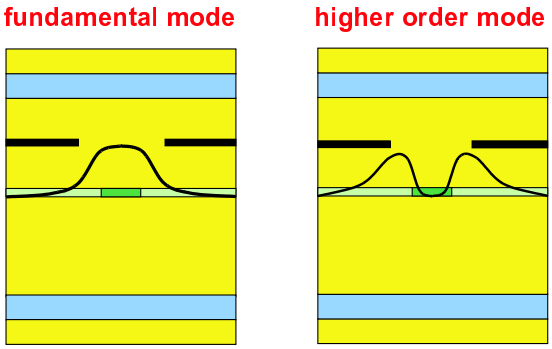

Figure 6.18: VCSEL combining both a proton-implantation and an oxide aperture to increase the modal stability.

6.3.5 Proton-implanted oxide-confined VCSELs

The effects of the gain diameter size relative to the oxide diameter size are exploited to increase modal stability in VCSELs which combine both a proton-implant and an oxide aperture [Seurin2000]. The diameter of the proton-implant aperture is chosen to be smaller than that of the oxide aperture, such that the effective gain diameter is also smaller than the oxide diameter (fig. 6.18). Such a device layout will significantly increase the threshold of the higher order modes (e.g. the TE_{01} mode, the lowest order mode with Bessel order 0 and TE polarisation) with respect to the fundamental mode. This is immediately clear from fig. 6.18, because for the higher order mode, almost all the gain is concentrated in regions where the optical field is very low. This high modal stability can lead to high-power devices which lase only in the fundamental mode, a desirable property for imaging or high-speed modulation.

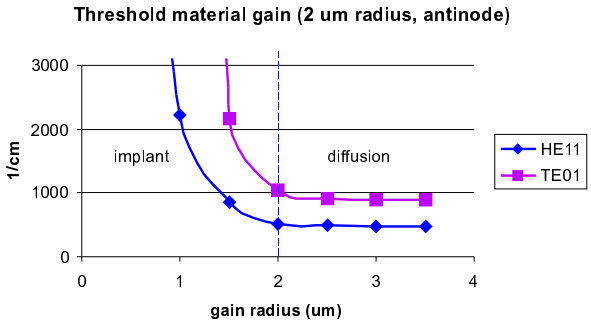

Figure 6.19: Threshold material gain for the HE_{11} and the TE_{01} mode as a function of gain diameter.

Fig. 6.19 shows the simulated behaviour of such a device, starting from the USC VCSEL structure with a 2 μm radius antinode oxide. The dashed line corresponds to a gain diameter equal to the oxide diameter, so to the right of this line we have situations where carrier diffusion increases the effective gain diameter. To the left of the line, the gain diameter is smaller than the oxide aperture, e.g. due to the presence of a proton-implantation with a smaller diameter. Obviously, decreasing the gain area leads to higher thresholds, both for the fundamental and the higher order mode. However, the relative threshold difference between these modes also becomes higher, leading to increased modal stability. Clearly, for practical applications there is a trade-off to be made between this increased modal stability and the higher threshold for the fundamental mode.

6.3.6 Tapered-oxide VCSELs

The oxidation rate of an AlGaAs layer is very strongly dependent on the Al fraction inside this layer. The higher the Al content, the faster the layer oxidises. For Al contents below 80%, the oxidation process becomes so slow that it is hardly noticeable. As an illustration of this dependence, we show the experimental data from fig. 6.20. This figure comes from another comparison which we coordinated within the framework of the COST268 action, this time with the aim of investigating how the oxidation process produces different results on different oxidation furnaces. For that purpose, two wafers with the same layer structure where cleaved into several samples and distributed to different European labs to perform oxidation experiments on. For a more detailed report on this comparison, we refer to [Bienstman2001b].

The exponential dependence of oxidation rate on Al content can be used to create oxide apertures which are tapered rather than abrupt (fig. 6.21). Instead of growing a single layer which will form the aperture

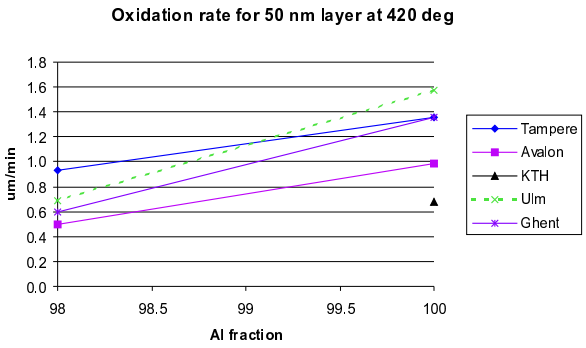

Figure 6.20: Oxidation rate of a 50 nm thick layer as a function of Al fraction.

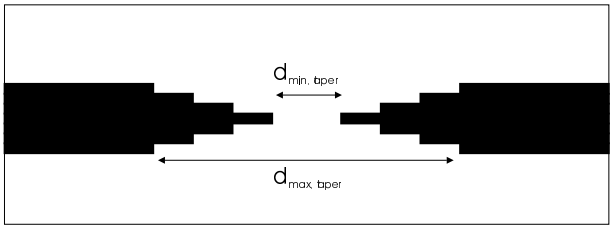

Figure 6.21: Tapered oxide aperture.

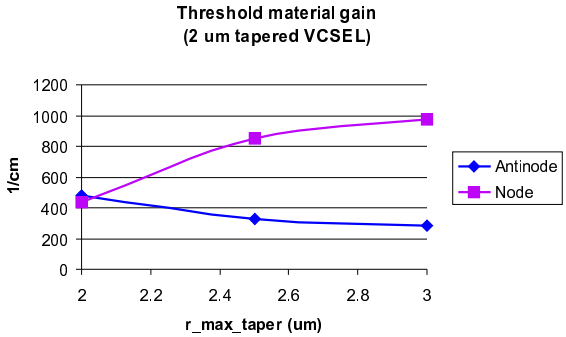

Figure 6.22: Threshold material gain for a tapered VCSEL.

after oxidation, many thinner layers are grown, each with a different Al content. If this structure is placed in the furnace for a fixed amount of time, the high Al content layers will have been oxidised deeper than the others. By carefully engineering the Al content profile, tapers with arbitrary profiles can be fabricated [Hegblom1999].

Compared to an abrupt oxide with equal thickness, a tapered oxide will provide lower confinement, but because the index change is much more gradual, scattering losses will be lower.

We illustrate this by replacing the abrupt thin oxides in the USC VCSEL with linearly tapered oxides. The inner taper radius is kept fixed at 2 μm, the outer radius is increased from 2 μm (no taper - abrupt oxide) to 3 μm. For the gain radius 3 μm is assumed. The simulation results are shown in fig. 6.22. For the node oxide, which already had a low confinement to start with, the even lower confinement caused by the tapering is so detrimental that the threshold goes up. Things are different for the antinode oxide, where the lower scattering outweighs the increased diffraction losses. Thresholds for the tapered antinode oxide are even lower than those of the abrupt oxide.

Clearly, this is only the tip of the iceberg, because careful control of the growth techniques and the oxidation process can provide much more sophisticated taper profiles than linear ones, like parabolic or asymmetric profiles. It is even hoped that the 'perfect' taper structure can be designed, which will act as a perfect lens that will focus the beam on a narrow spot inside the active region, without any scattering or diffraction losses [Noble2000]. Such devices could potentially have extremely low threshold currents. This is a promising area for further research. However, engineering the perfect VCSEL cavity will not be trivial because at this level of sophistication, the precise carrier diffusion profiles will have to be taken into account, in addition to other effects like thermally induced refractive index changes.

6.4 Conclusion

We have seen that designing a confinement structure in VCSELs involves a delicate balance between diffraction and scattering losses, between the conflicting demands of high confinement inside the confining structure and small refractive index steps to reduce scattering for light entering the confining structure. Moreover, these issues cannot be considered in isolation from the rest of the VCSEL design: other aspects like cavity length or carrier diffusion have to be taken into account. All of this makes it hard to predict the optical performance of a particular design intuitively, which is precisely why we need optical simulation tools that are both rigorously accurate and computationally efficient. We have shown that CAMFR is very well suited to fulfill such a role, even more so when it will be coupled in the future to e.g. thermal and carrier models, to give a full picture of all the physical phenomena at work in these devices.

Chapter 7

Resonant-Cavity LEDs

> "Models are to be used, not believed."
> H. Theil, "Principles of Econometrics"

In the previous chapter we studied laser modes, which are field profiles that can be sustained indefinitely without the presence of external sources. In this chapter, we will investigate what happens when there is a source present inside the cavity. First, we will show how to use eigenmode expansion to calculate the response of a cavity to an arbitrary source. Then, we will turn our attention to a specific kind of source, the dipole current source. We will use this source to study spontaneous emission in microcavities, which is an important aspect in resonant-cavity light-emitting diodes. After discussing these RCLEDs, we will also propose a novel type of RCLED that allows more directive radiation patterns.

7.1 Source inside a cavity

Consider a cavity which we divide at an arbitrary position in a top and a bottom part, just like we did in the previous chapter. In the plane separating these two parts, we place a source that radiates with a specific field pattern. The exact nature of this field pattern is arbitrary, we only need to know its expansion coefficients $\mathbf{A}_0$ in terms of the eigenmodes of the layer containing the source. In general, the source will radiate both upwards and downwards, so it is fully characterised by the two expansion vectors $\mathbf{A}_{up,0}$ and $\mathbf{A}_{do,0}$.

We assume that these two vectors describe the source field as it would be in a z-invariant layer, i.e. in the absense of any reflecting structures in the z-direction. The presence of the cavity will cause reflections, and this interference will give rise to different total field profiles $\mathbf{A}_{up}$ and $\mathbf{A}_{do}$, which we will now calculate. Let us concentrate for

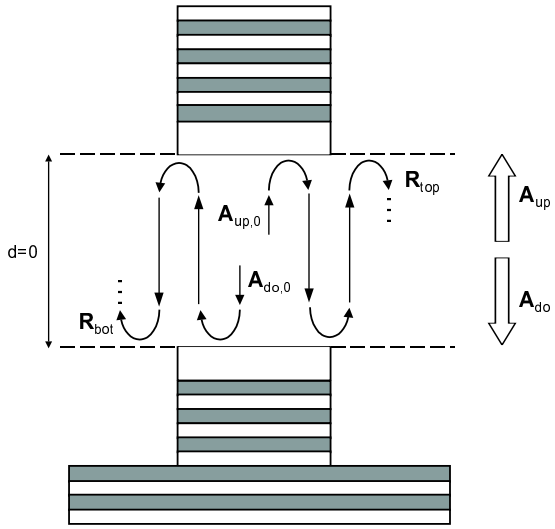

Figure 7.1: Cavity containing a source.

the moment on the upward field $\mathbf{A}_{up,0}$. This field distribution will keep on bouncing back and forth between the top and bottom mirror, such that the equilibrium upward field distribution just above the source will be

$$\mathbf{A}_{up,0} + \mathbf{R}_{bot} \cdot \mathbf{R}_{top} \cdot \mathbf{A}_{up,0} + \mathbf{R}_{bot} \cdot \mathbf{R}_{top} \cdot \mathbf{R}_{bot} \cdot \mathbf{R}_{top} \cdot \mathbf{A}_{up,0} + \ldots \tag{7.1}$$

We can write this as

$$\sum_{i=0}^{\infty} \left(\mathbf{R}_{bot} \cdot \mathbf{R}_{top}\right)^i \cdot \mathbf{A}_{up,0} \tag{7.2}$$

This sum converges to[1]

$$\left(\mathbf{I} - \mathbf{R}_{bot} \cdot \mathbf{R}_{top}\right)^{-1} \cdot \mathbf{A}_{up,0} \tag{7.3}$$

We must not forget that also the downward radiated field contributes to the total upward field. After reflecting at the bottom mirror, it goes through the same sequence of reflections as in eq. 7.1. Therefore, the total upward equilibrium field can be written as

$$\mathbf{A}_{up} = \left(\mathbf{I} - \mathbf{R}_{bot} \cdot \mathbf{R}_{top}\right)^{-1} \cdot \left(\mathbf{A}_{up,0} + \mathbf{R}_{bot} \cdot \mathbf{A}_{do,0}\right) \tag{7.4}$$

Similarly, the total downward field just below the source is

[1] This can be seen easily by subtracting $\mathbf{I}+\mathbf{Q}+\mathbf{Q}^2+..=\mathbf{X}$ and $\mathbf{Q}+\mathbf{Q}^2+\mathbf{Q}^3+...=\mathbf{Q}\cdot\mathbf{X}$ and solving the resulting expression for $\mathbf{X}$.

$$\mathbf{A}_{do} = \left(\mathbf{I} - \mathbf{R}_{top} \cdot \mathbf{R}_{bot}\right)^{-1} \cdot \left(\mathbf{A}_{do,0} + \mathbf{R}_{top} \cdot \mathbf{A}_{up,0}\right) \tag{7.5}$$

Using these field distributions as excitations for respectively the top and bottom half of the cavity, we can calculate the field at an arbitrary position, as per section 2.5. This completely characterises the response of the cavity to the source.

Note that the presence of the source introduces discontinuities in the field profile. E.g., the upward field $\mathbf{A}_{up}$ just above the source is different from the upward field just below the source, which is equal to $\mathbf{R}_{bot} \cdot \mathbf{A}_{do}$.

It is interesting to point out that eq. 7.4 and 7.5 become singular when the product of $\mathbf{R}_{bot}$ and $\mathbf{R}_{top}$ equals the unity matrix. This is not surprising, because this is exactly the lasing resonance condition from the previous chapter. The fact that these factors become infinite, means that even for zero $\mathbf{A}_{up,0}$ and $\mathbf{A}_{do,0}$ the total fields $\mathbf{A}_{up}$ and $\mathbf{A}_{do}$ can be finite. In other words, a laser mode can exist without being sustained by external sources.

7.2 Eigenmode expansion of a dipole current source

In this section, we will calculate the field expansion coefficients $\mathbf{A}_{up,0}$ and $\mathbf{A}_{do,0}$ of an arbitrary current distribution as a source. This will be easily applied to a point dipole current source. We can then use the techniques outlined in the previous section to calculate the response to such a source, which gives us nothing other than the Green's function of the cavity.

Consider a z-invariant layer containing a current distribution $\mathbf{J} = \mathbf{J}(x,y)\delta(z)$ at $z = 0$. We want to calculate the expansion of the fields excited by this source in terms of the eigenmodes of the z-invariant layer. For that, we make use of the Lorentz reciprocity theorem (eq. 2.15) [Derudder2001a]:

$$\int\int_S (\mathbf{E}_1 \times \mathbf{H}_2 - \mathbf{E}_2 \times \mathbf{H}_1) \cdot d\mathbf{S} = \int\int\int_V (\mathbf{J}_1 \cdot \mathbf{E}_2 - \mathbf{J}_2 \cdot \mathbf{E}_1) dV \tag{7.6}$$

For the quantities with subscript 1, we choose the fields excited by the current distribution $\mathbf{J}_1 = \mathbf{J}$, and we expand them in terms of the eigenmodes:

$$\mathbf{E}_1 = \begin{cases} \sum A_i^+ \mathbf{E}_i^+ \ , \ z > 0 \\ \sum A_i^- \mathbf{E}_i^- \ , \ z < 0 \end{cases} \tag{7.7}$$

$$\mathbf{H}_1 = \begin{cases} \sum A_i^+ \mathbf{H}_i^+ \ , \ z > 0 \\ \sum A_i^- \mathbf{H}_i^- \ , \ z < 0 \end{cases} \tag{7.8}$$

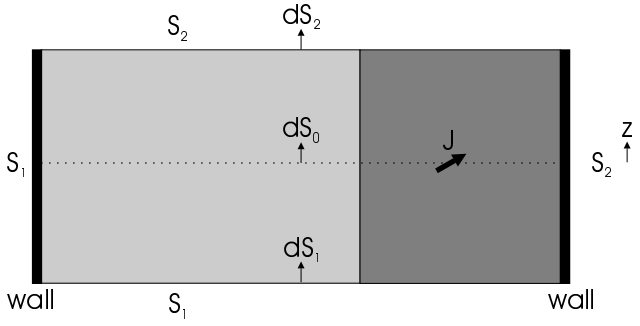

Figure 7.2: Volume for the evaluation of the Lorentz reciprocity theorem.

The $+(-)$ superscript denotes eigenmodes propagating in the positive (negative) z-direction. Note that we take into account the presence of a discontinuity caused by the source.

For the quantities with subscript 2, we choose a sourceless situation ($\mathbf{J}_2 = 0$), with the m-th eigenmode as field profile. This is possible, precisely because eigenmodes can exist without the need for external sources.

$$\mathbf{E}_2 = \mathbf{E}_m^+ \tag{7.9}$$
$$\mathbf{H}_2 = \mathbf{H}_m^+ \tag{7.10}$$

For the evalution of the integrals, we take the volume V bounded by the surface $S = S_1 + \delta S_1 + S_2 + \delta S_2$ as indicated in fig. 7.2. Because of the presence of $\delta(z)$ in $\mathbf{J}$, the volume integral on the RHS of eq. 7.6 reduces to a surface integral over S_0 at $z = 0$.

For the LHS we already know that the contribution from $\delta S_1 + \delta S_2$ will vanish in most cases due to the nature of the boundary conditions (see section 3.5.1). We can then write

$$-\int\int_{S_1} (\mathbf{E}_1 \times \mathbf{H}_2 - \mathbf{E}_2 \times \mathbf{H}_1) \cdot d\mathbf{S}_1 +$$
$$\int\int_{S_2} (\mathbf{E}_1 \times \mathbf{H}_2 - \mathbf{E}_2 \times \mathbf{H}_1) \cdot d\mathbf{S}_2 = \int\int_{S_0} \mathbf{J} \cdot \mathbf{E}_2 dS_0 \tag{7.11}$$

Note from fig. 7.2 that all the $d\mathbf{S}$ vectors are oriented along $+z$.

We know that eigenmodes of z-invariant layers are orthogonal. We also assume that their amplitudes are suitable chosen such that

$$\int\int_S \left(\mathbf{E}_{i,t}^+ \times \mathbf{H}_{m,t}^+\right) \cdot d\mathbf{S} = \delta_{im} \tag{7.12}$$

Only the tangential field components contribute to this vectorial product. Thanks to this orthonormality, the summations that make up $\mathbf{E}_1$ and $\mathbf{H}_1$ in eq. 7.11 will reduce to the terms with index m:

$$-A_m^- \int\int_{S_1} \left(\mathbf{E}_{m,t}^- \times \mathbf{H}_{m,t}^+ - \mathbf{E}_{m,t}^+ \times \mathbf{H}_{m,t}^-\right) \cdot d\mathbf{S}_1 +$$
$$A_m^+ \int\int_{S_2} \left(\mathbf{E}_{m,t}^+ \times \mathbf{H}_{m,t}^+ - \mathbf{E}_{m,t}^+ \times \mathbf{H}_{m,t}^+\right) \cdot d\mathbf{S}_2 = \int\int_{S_0} \mathbf{J} \cdot \mathbf{E}_m^+ dS_0 \quad (7.13)$$

From eq. 2.24 follows the relation between the field profiles of the backward eigenmode and those of the forward one: $\mathbf{E}_{m,t}^- = \mathbf{E}_{m,t}^+$ and $\mathbf{H}_{m,t}^- = -\mathbf{H}_{m,t}^+$.

All of this allows us to simplify eq. 7.11 as

$$-2A_m^- = \int\int_{S_0} \mathbf{J} \cdot \mathbf{E}_m^+ dS_0 \quad (7.14)$$

Similarly, if we use for $\mathbf{E}_2$ and $\mathbf{H}_2$ the backward eigenmode instead of the forward one, we arrive at

$$-2A_m^+ = \int\int_{S_0} \mathbf{J} \cdot \mathbf{E}_m^- dS_0 \quad (7.15)$$

These two expressions completely describe the expansion of an arbitrary current source in terms of eigenmodes.

For the special case of a dipole point current source at $\mathbf{r} = \mathbf{r}_0$, $\mathbf{J}$ is simply $\mathbf{J}_0 \delta(\mathbf{r} - \mathbf{r}_0)$, and we get

$$-2A_m^- = \mathbf{J}_0 \cdot \mathbf{E}_m^+(\mathbf{r}_0) \quad (7.16)$$
$$-2A_m^+ = \mathbf{J}_0 \cdot \mathbf{E}_m^-(\mathbf{r}_0) \quad (7.17)$$

This means that the excitation strength of a particular eigenmode is proportional to its field amplitude at the source. E.g., if the source is placed at a null of an eigenmode's electrical field, this mode will not be excited. The same holds for a dipole oriented perpendicular to the electric field. All this is also very acceptable from a physical point of view.

A few concluding remarks to end this section. The discussion so far was completely independent of the coordinate system used. Once again this is very beneficial, because the implementation in CAMFR can proceed using abstract waveguides. Secondly, the description of the field radiated by a dipole is also more general than some incarnations of the plane-wave method found in literature (see e.g. [Lukosz1977], [De Neve1997a] and [Benisty1998]), which are only valid inside a uniform lossless medium. Using the approach presented here, we can easiliy accomodate laterally varying media and lossy materials.

7.3 Modification of spontaneous emission

A dipole current source is a good model for the radiation emitted by the recombination of an electron and a hole [Loehr1998][Demeulenaere1997]. It is therefore useful if we want to model the spontaneous emission process in semiconductor light emitters. This spontaneous emission is not an intrisic property of the active layer, but it can be modified by the optical environment in which the source is placed. In the parlance of this chapter, placing a dipole with intrinsic radiation profile $(\mathbf{A}_{up,0}, \mathbf{A}_{do,0})$ inside a cavity will modify its radiation pattern to $(\mathbf{A}_{up}, \mathbf{A}_{do})$. This is not necessarily a power-conserving operation. In some cases, the dipole will radiate less power than when placed in vacuum, in other cases more. This is known as inhibited or enhanced spontaneous emission.

Inhibited spontaneous emission occurs e.g. when the dipole is placed inside a cavity that does not support any modes the dipole can couple to. In this case, the light effectively 'has nowhere to go', so spontaneous emission cannot occur and the dipole stays in the excited state indefinitely. On the other hand, if the cavity has a mode that is resonant with the dipole radiation, the emitted power will be resonantly enhanced. These resonant modes are very efficient to carry energy away from the dipole, so it will decay faster than in free space and we have *enhanced spontaneous emission.*

In this context, we can define a dipole decay rate γ, which we relate to the emitted dipole power P as follows:

$$\frac{\gamma}{\gamma_{vac}} = \frac{P}{P_{vac}} \tag{7.18}$$

The subscript vac refers to the situation in vacuum, the quantities without subscript to a dipole placed in a cavity. Obviously, if the dipole radiates more power, it is because it can couple well to the cavity modes, which gives rise to a shorter life-time and hence an increased decay rate.

A well-known example to illustrate these effects is to study the spontaneous emission of a dipole placed halfway between two perfectly conducting metal mirrors. This point dipole is oriented parallel to the metal plates and the medium between the mirrors is air (fig. 7.3). Because of the rotational symmetry, the angular orientation of this horizontal dipole is irrelevant. This is a 3D problem, in the sense that radiation is allowed to travel outside of the plane of the figure. The modification of spontaneous emission rate can easily be calculated from eq. 7.18. We know from the previous sections how to calculate the field profiles inside the cavity, and the total power emitted by the dipole follows then quite simply from integrating the flux of the Poynting vector through the surfaces S_1 and S_2.

For the simple geometry of fig. 7.3, it is possible to come up with a closed-form solution for the dipole decay rate [Abram1998]. However,

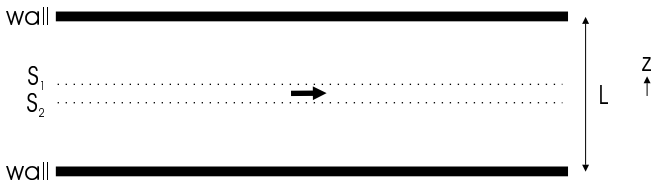

Figure 7.3: Dipole radiating between two metal plates.

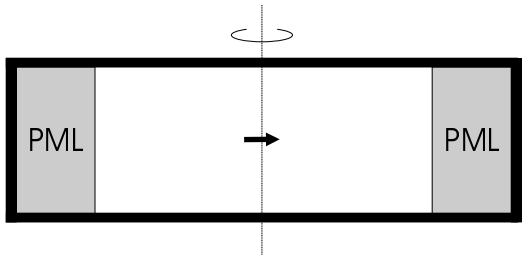

Figure 7.4: Geometry of fig. 7.3 transformed to a circular closed geometry.

to validate the techniques outlined in this chapter, we also calculated the decay rate numerically with CAMFR, using a cylinderical geometry to model the 3D problem. We close off the domain in the ρ-direction with hard walls clad with PML (fig. 7.4).

In theory, we should consider circular modes with any Bessel order for the calculations. However, for this particular geometry, it turns out that the dipole only excites modes with Bessel order 1. Fig. 7.5 shows the calculated spontaneous emission rate of the dipole as a function of the distance L between the mirrors. The results are normalised to the dipole emission rate in vacuum. The numerical and the analytic results coincide almost perfectly, apart from some oscillations at the discontinuities.

Several observations can be made from fig. 7.5. For $L < \lambda/2$, the dipole does not radiate at all, because all the modes in the cavity are below cut-off. At $L = \lambda/2$, we have maximum spontaneous emission, a factor 3 enhancement over the vacuum rate. As the cavity length increases, the effects of enhancement and inhibition become less pronounced, and the rate approaches the vacuum rate.

It is also interesting to take a look at the convergence behaviour of these calculations as a function of the number of modes N retained in the eigenmode expansion. Fig. 7.6 show such results, for a cavity length $L/\lambda = 0.75$ and for different radii of the 'pill box' from fig. 7.4 (r =15, 25 and 35 μm). The imaginary radius providing PML absorption is -0.1j. After some initial transitory behaviour, all curves con-

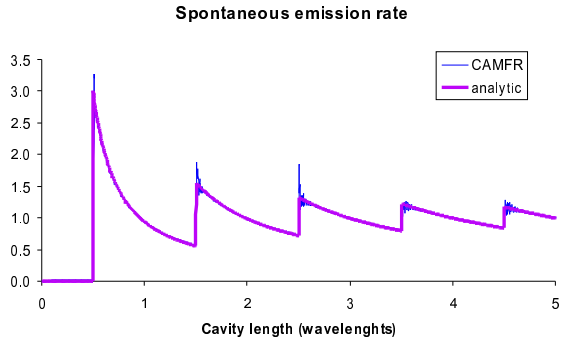

Figure 7.5: Spontaneous emission rate for a dipole between two mirrors.

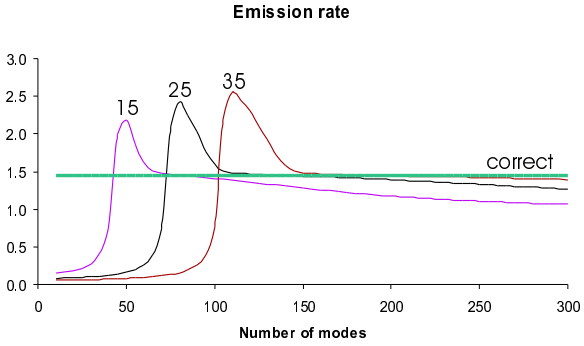

Figure 7.6: Convergence properties of the emission rate.

verge to a plateau very close to the correct value of 1.444, as calculated analytically. However, when increasing the number of modes even further, the curves drop off and approach the wrong value of 1. It also turns out that when we increase the absorption in the PML, the width of these plateaus of correctness decreases, until they finally disappear altogether for very high absorption levels.

The explanation for this behaviour is most likely that the set of PML modes is mathematically not complete, as we already suggested in section 3.4.4. When considering high absorption in the PML (or moderate absorption and high order modes), the PML modes converge to the leaky modes, which are known not to form a complete set. For practical purposes however, this rarely turns out to be a problem. When in doubt, a convergence analysis like the one in fig. 7.6 can be performed and the presence of 'plateaus of correctness' indicates convergence of the physical quantities calculated.

7.4 The RCLED

Placing a current source inside a cavity does not only alter the spontaneous emission rate, it also affects the radiation profile. The latter effect is much more important for practical applications, because it helps to increase the outcoupling efficiency of light-emitting diodes (LEDs).

Traditional LEDs consist of a layer of high index material placed in air. Not all waves emitted by a source inside the material will be able to reach the surrounding air region. Only the radiation that travels perpendicular to the semiconductor-air interface will be coupled out. Waves that travel too obliquely will undergo total internal reflection at the interface and will be sent back into the semiconductor. For GaAs with a refractive index of 3.5, calculations show that only 2% of the light emitted in the active region will be able to reach the outside world.

Resonant-cavity LEDs (RCLEDs) overcome this problem by placing the active layer inside a cavity. The cavity is engineered such that radiation perpendicular to the interface is resonantly enhanced by constructive interference (fig. 7.7). Similarly, oblique radiation is suppressed due to destructive interference. Because of this internal spatial redistribution of the radiation profile, a larger fraction of the light will be travelling perpendicular to the interface, leading to a lower loss from total internal reflection and to higher outcoupling efficiencies. This concept was first introduced in [Schubert1994], and later perfected by [De Neve1997b] and [Wierer1999], leading to devices with extraction efficiencies of over 20%.

To illustrate the physical principles behind the RCLED, a representation in k-space is very instructive [Brorson1990]. In this qualitative diagram, the spontaneous emission profile is represented by the surface of a sphere. What is meant here is the *intrinsic* spontanous

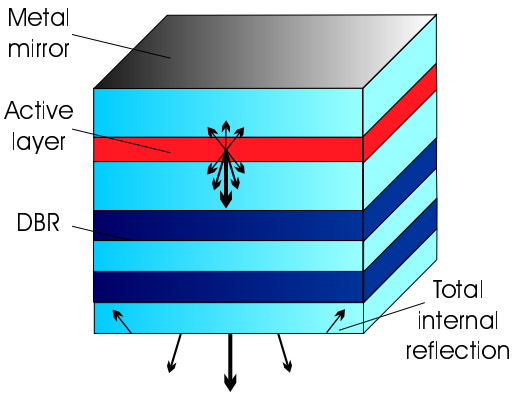

Figure 7.7: Resonant-cavity LED.

emission profile, i.e. in the absence of any cavity. This emission is assumed to be isotropic, i.e. equally strong in each direction (hence the sphere). In the the case of a monochromatic (single wavelength) emitter, the thickness of this sphere surface is delta-like. For more realistic sources, the emission takes place in a wavelength range around the central wavelength, so the sphere surface has a certain thickness.

When the emitter is placed inside a cavity, radiation in certain emission angles will be resonantly enhanced due to constructive interference. The condition for this constructive interference is a round-trip phase change of an integer multiple of 2π:

$$\phi_{top}(\theta) + \phi_{bot}(\theta) - 2Lk\cos\theta = l2\pi \ , \ l \text{ integer} \tag{7.19}$$

Here, ϕ_{top} and ϕ_{bot} are the phases of the reflection coefficients of the top and the bottom mirror. The cavity has a length of L and the angle of propagation inside the cavity is given by θ. The length of the wavevector in the material is k. From eq. 7.19 it is trivial to derive the k-vectors that satisfy the resonance condition:

$$k_{res}(\theta) = \frac{\phi_{top}(\theta) + \phi_{bot}(\theta) - l2\pi}{2L\cos\theta} \ , \ l \text{ integer} \tag{7.20}$$

Eq. 7.20 can easily be represented in the angular k-diagram as a certain surface[2]. Because this resonant enhancement has a non-zero bandwidth, we draw this surface with a certain thickness.

All of this eventually leads to a representation like the one in fig. 7.8. The interpretation of this figure is as follows. Points in k-space where there is a significant overlap between the spontaneous emission sphere k_{sp} and the cavity resonance surface k_{res} correspond to radiation emitted by the active layer that is enhanced by the cavity. Similarly, points

[2] Or more precisely a set of surfaces for each value of l.

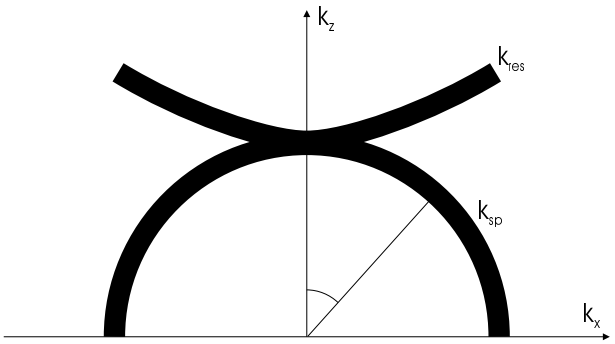

Figure 7.8: k-space diagram for a resonant cavity LED.

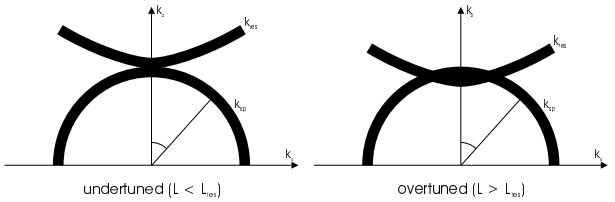

Figure 7.9: k-diagram for an undertuned and an overtuned cavity.

with no overlap correspond to inhibited spontaneous emission. From such a representation, it is immediately clear that the cavity changes the emission profile of the active layer from an isotropic one to a profile where radiation is emitted predominantly along $\theta = 0$, i.e. perpendicular to the mirrors.

The $k-$space diagram is also very useful to understand what happens when we change the cavity length L (fig. 7.9). Decreasing the cavity length L will cause the resonance surface to shift upwards (eq. 7.20). For these so-called undertuned cavities, the total overlap between the sphere and the surface is a lot smaller and is concentrated in a narrower angular range. These kinds of cavities will therefore have narrower radiation profiles and lower extraction efficiencies. For overtuned devices with too long cavity lengths, the opposite is true. The resonance surface shifts downwards, so that the intersection between the sphere and the surface is no longer at $\theta = 0$, but off-axis. In this case, the radiation profile will be much wider. For very high overtuning, it will even have sidelobes. Additionally, the total overlap between the sphere and the surface will be larger, so the extraction efficiency increases with respect to the resonant case.

To illustrate these points, we performed some CAMFR simulations on a very simple planar RCLED device structure (table 7.1), that we simulated using a cylindrical geometry as per fig. 7.4. It is a bottom-

	Thickness (nm)	Material	Refractive index
air		air	1.0
top mirror	120	Au	0.2-6.5j
	45.5	GaAs	3.5
active layer	5	QW	3.5-0.1j
	d	GaAs	3.5
bottom mirror	158	AlOx	1.55
substrate		GaAs	3.5
exit medium		air	1.0

Table 7.1: Layer structure of RCLED with oxidised bottom DBR mirror.

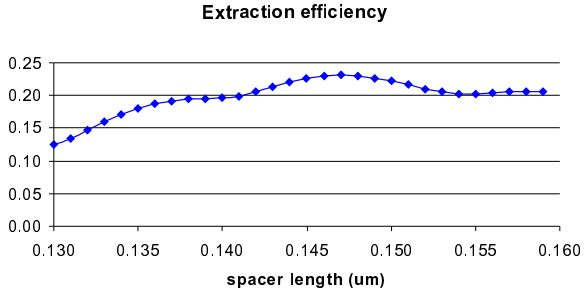

Figure 7.10: Extraction efficiency η_{ex} of the RCLED from table 7.1.

emitting structure with an operating wavelength of 980 nm. The top mirror is made of metal and is also used as a contact to inject the current. For simplicity, we assume that the QW only emits a single wavelength. We assume a dipole oriented parallel to the layer structure. The bottom mirror is a single-layer DBR consisting of oxidised material. Only one layer is used, because multilayer DBRs would quickly yield very high reflectivities because of the high index contrast. In RCLEDs, it is better to keep the mirror reflectivity moderate. Otherwise, the bandwidth of resonant enhancement would become very narrow, and only a small fraction of the emitted light would be resonantly enhanced. Because the vast majority of the light would not benefit from the microcavity effect, the extraction efficiency of such a device would be quite low.

Fig. 7.10 shows the calculated extraction efficiency η_{ex} of the device from table 7.1, as a function of the spacer thickness d. This quantity η_{ex} can be calculated by dividing the downward power flux in the bottom air exit medium and the total power flux leaving the active layer:

$$\eta_{ex} = \frac{P_{air,down}}{P_{QW,up} + P_{QW,down}} \tag{7.21}$$

For the simulations, the lower substrate-air interface is assumed to

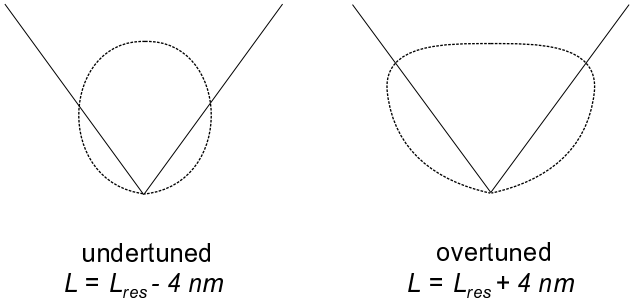

Figure 7.11: Far field profiles in RCLEDs.

reflect incoherently, i.e. the reflections there do not contribute to the microcavity effect. This is a very valid assumption, because the substrate thickness of a few hundred μm is a lot longer than the coherence length of the light emitted by a typical QW.

From fig. 7.10, it can be seen that the maximum extraction efficiency occurs at around 148 nm spacer thickness. The resonance condition is at 140 nm, so clearly longer cavities can increase the extraction efficiency.

The far-field radiation profiles for these devices are shown in fig. 7.11, both for an undertuned and an overtuned device. For short cavities, the radiation profile is narrower, as we already indicated in the discussion of fig. 7.9.

Not all the power in this far field profile will be useful for practical applications. E.g., if we want to couple the emitted light into a plastic optical fibre without any lenses, only the power propagating within a certain acceptance cone (the straight lines in fig. 7.11) will be able to enter the fibre. In this context, we define the sine of half the top angle of this cone as its numerical aperture (NA). For typical plastic optical fibres, this NA is about 0.5. Fig. 7.12 shows the fraction of the optical power in air that lies within a NA of 0.5:

$$\eta_{NA} = \frac{P_{air,down,NA}}{P_{air,down}} \tag{7.22}$$

As expected from the field profiles, η_{NA} will be a lot smaller for overtuned cavities.

Both undertuned and overtuned cavities are valid design options. Overtuned cavities are used when one wants to maximise the total output power, at the expense of a lower directivity of the beam profile [De Neve1997b]. Undertuned cavities are ideally suited for applications where a narrow beam profile is required [Bockstaele1999]. The price to pay for this directivity is a lower output power.

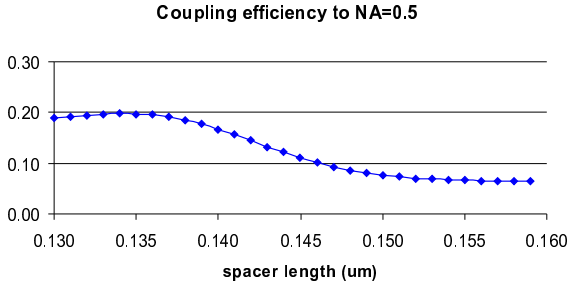

Figure 7.12: Fraction of the outcoupled power within NA=0.5.

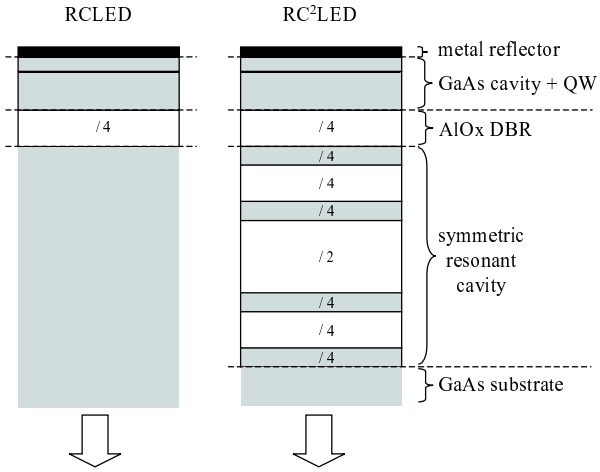

Figure 7.13: The RCLED vs. the RC^{2}LED.

7.5 The RC^{2}LED

In this section, we propose a novel RCLED design that also allows to increase the directivity of the beam profile, and that has additional benefits as compared to undertuning the cavity. This so-called RC^{2}LED design was published in [Bienstman2000] and was the subject of a patent application.

The idea behind this approach is to use a different structure for the bottom outcoupling mirror (fig. 7.13). This bottom mirror consists of two parts. A first part is a traditional DBR mirror, just like in the RCLED. A second part placed below the first one is a symmetric cavity that has the same resonance wavelength as the cavity containing the active layer. It is also made of the same material as the cavity, namely AlOx and GaAs. It is easy to show that this second cavity is completely transparent for the resonance wavelength under normal

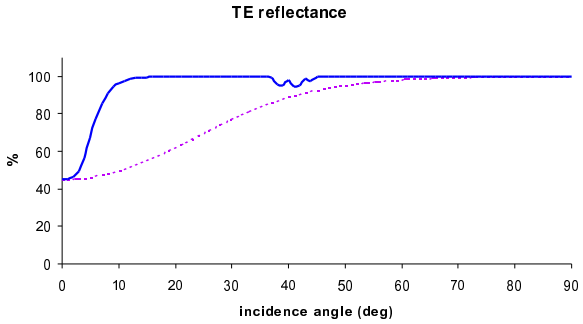

Figure 7.14: TE reflectance for a DBR (dashed line) and a DBR+RC (solid line), as seen from the GaAs cavity.

incidence. Therefore, the normal incidence reflectivity of the bottom mirror is only determined by its first part. This decoupling is important, because it allows us to choose the bottom reflectivity low enough. As already mentioned in the previous section, this is vital if we want to design efficient devices.

It is well-known that the reflection of a resonant cavity is always very high, except when the incident light corresponds to a cavity resonance. Thus, the combination of a traditional DBR with a symmetric resonant cavity (RC) yields a structure where the off-axis reflectivity is always very high, while the normal incidence reflectivity is set by the DBR. This is illustrated in fig. 7.14, comparing the reflectivity at 980 nm for the DBR and the DBR + RC. From this figure, it is already apparent that the RC^{2}LED will have a narrow radiation profile, because the transmission of the bottom mirror is only significant in a much smaller cone as compared to the RCLED.

This could lead to the impression that the outcoupling efficiency of the RC^{2}LED will also be a lot lower, because only a narrow cone of the emitted radiation will be able to reach the outside world, while the rest is effectively thrown away. However, this effect is compensated to a large degree by the peculiar phase properties of the outcoupling mirror, as shown in fig. 7.15. Provided the DBRs in the RC have an odd number of layers and provided the refractive index contrast is high enough, the combination of a DBR and a RC has a significant *negative* angular penetration depth, i.e. a negative phase slope. This means that instead of a single resonance at normal incidence, the RC^{2}LED will have an additional off-axis resonance. This leads to a higher overlap between the spontaneous emission sphere and the resonance surface in the k-diagram (fig. 7.16).

We illustrate this behaviour with some simulation results on the same structure from table 7.1, but with an extra RC as indicated in fig. 7.13. The field profiles of the RC^{2}LED are illustrated in fig. 7.17,

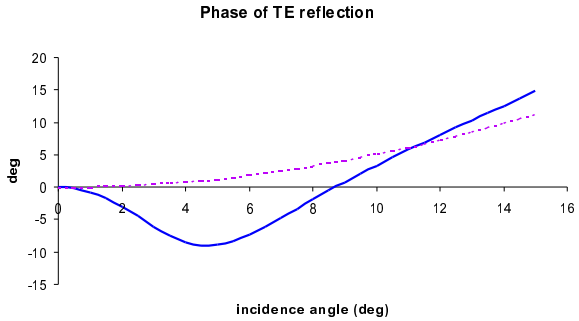

Figure 7.15: Phase of TE reflection for a DBR (dashed line) and a DBR+RC (solid line).

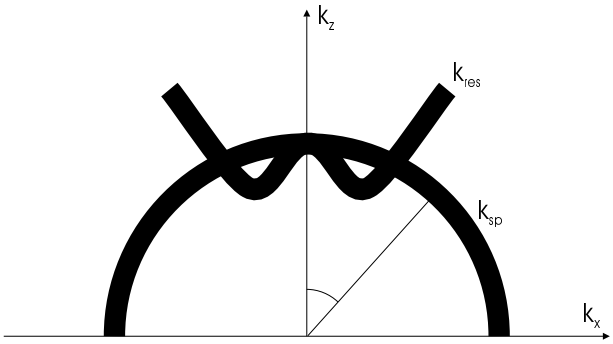

Figure 7.16: k-space diagram of an RC^{2}LED.

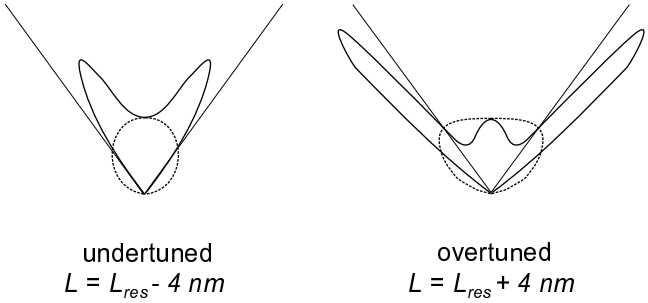

Figure 7.17: Far field profiles of the RCLED (dashed line) and the RC^{2}LED (full line).

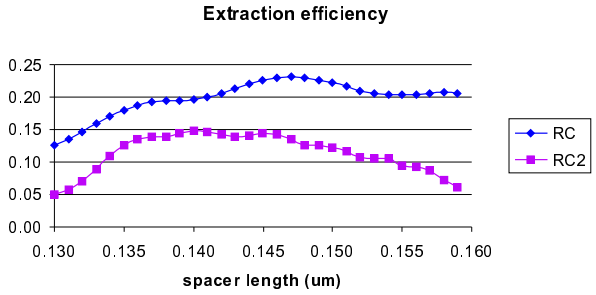

Figure 7.18: Extraction efficiency η_{ex} of the RCLED (dashed line) and the RC^{2}LED (full line).

both for an undertuned and an overtuned device. For normal incidence, the RC is transparent and the emission of the RC^{2}LED is the same as that of the RCLED. However, the RC^{2}LED has additional sidelobes corresponding to the extra resonance in the k-diagram. As the cavity length increases, the angle corresponding to these sidelobes increases, consistent with a resonance surface shifting down in the k-diagram.

Although the extraction efficiency of the RC^{2}LED is slighty lower than in the case of the RCLED (fig. 7.18), the power coupled to a limited NA is significantly higher (fig. 7.19).

When increasing the cavity length, η_{NA} suddently drops, which can be explained by the sidelobes shifting outside of the acceptance cone of the fibre in fig. 7.17.

Ultimately, the relevant figure of merit for these devices is the product of η_{ex} and η_{NA}, the fraction of the light emitted by the active layer that makes it to the fibre. From fig. 7.20, we can see that the RC^{2}LED offers a two-fold improvement over a classical RCLED design, which can be quite an important asset when creating optical links, where the

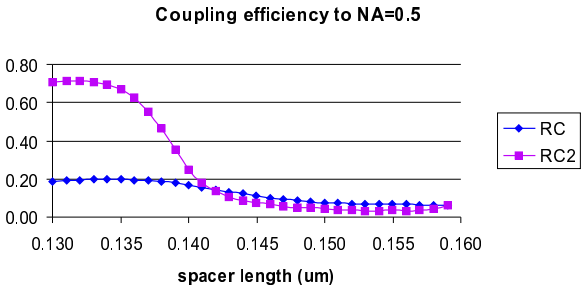

Figure 7.19: Fraction η_{NA} in NA=0.5 of the RCLED (dashed line) and the RC^{2}LED (full line).

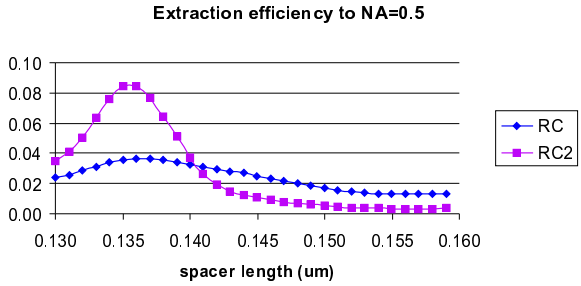

Figure 7.20: Extraction efficiency $\eta_{ex} \cdot \eta_{NA}$ to NA=0.5 of the RCLED (dashed line) and the RC^{2}LED (full line).

power budget is always a critical issue.

Currently, efforts are underway to fabricate this device and test these theoretical predictions.

Chapter 8

Future perspectives

"It's tough to make predictions, especially about the future."
Yogi Berra

8.1 Conclusions

During this work, we extended the eigenmode expansion method to use advanced boundary conditions, like perfectly matched layers or transparent boundary conditions. This allows the method to correctly model structures with radiation losses, whereas before, this radiation would reflect off the hard boundary walls and would disturb the results. Not only does this increase the accuracy of the results, it also improves the speed of the method because the boundaries can now be placed much closer to the structure under study. Due to this reduction in computational volume, a smaller number of modes N is required to achieve convergence, which brings about important speed-ups because the calculation time scales as N^3.

We implemented these numerical techniques in a flexible, easy to use CAD framework called CAMFR, short for CAvity Modelling FRamework. Through the use of object-oriented design techniques, the implementation of the models can proceed without explicit knowledge of the details of the underlying spatial geometry, which makes it very easy to extend CAMFR with new geometries.

Currently, two geometries are implemented in CAMFR:

- a 2D cartesian geometry, with an arbitrary number of index steps in the z-direction and in the lateral direction.
- a 3D cylindrical symmetric geometry, with an arbitrary number of index steps in the z-direction and per layer a single index step in the radial direction.

For these geometries (and also for future geometries), the following functionality is available:

- fully vectorial field profiles inside a structure, including the reflection and transmission coefficients for arbitrary incident fields.
- band diagrams of an infinite periodic repetition of a basic structure.
- lasing modes of an arbitrary cavity, including resonance wavelength, threshold material gain and field profiles.
- response of an arbitrary cavity to an arbitrary current source.

The calculation times of CAMFR compare extremely favourably to other methods, especially to those based on spatial discretisation, like e.g. FDTD. This is largely due to the fact that eigenmode expansion is able to treat regions of constant refractive index analytically. CAMFR calculation times are typically on the order of seconds or minutes, instead of hours or days using other methods.

We have shown that CAMFR can be used to model a large variety of optical structures and devices. We applied it to study band diagrams in photonic crystals and to calculate the properties of waveguides, bends and splitters in these photonic crystals. We performed a detailed design study on different optical confining mechanisms in vertical-cavity surface-emitting lasers and elucidated some of the trade-offs to make in these matters. Finally, we investigated the process of spontaneous emission in microcavities and proposed a novel type of resonant-cavity light-emitting diode that has an increased extraction efficiency to optical fibre.

8.2 Future work

There are several topics that could be the subject of further work and extensions to CAMFR.

The most obvious one is without doubt the extension to other geometries. One particularly attractive geometry, but also a very challenging one, is the extension to 3D cartesian structures. While this extension is conceptually simple, it is numerically far from straightforward, among others because of the larger number of modes that will be required to model 3D structures. New techniques will therefore be required in order to keep the computational load within reasonable bounds. A possible route consists in exploiting the sparsity of the matrices involved. Initial investigations with 2D structures show that even there, many elements in the matrices are close to zero, so incorporating sparse matrix solvers in CAMFR will provide additional speed-ups.

We hope that CAMFR could eventually become a valuable option for full-vectorial 3D simulations, which are currently very challenging to perform and demand considerable investments in computer hardware. It is our belief that eigenmode expansion could offer significant advantages over FDTD, especially for structures with index profiles consisting of large piecewise constant sections.

Slightly less ambitious, but nevertheless very interesting from a practical point of view, would be the extension to cylindrical waveguides with an arbitrary number of radial index steps. For VCSELs, this would allow CAMFR to model more realistic gain profiles inside the QW, by doing a staircase approximation. This is important if we want to couple our optical model to other thermal or electrical models. Secondly, this geometry could be used to study resonators with cylindrical leek-like Bragg reflectors [Labilloy1998], which represent a possible route to full 3D optical confinement inside a cavity. The modelling of spontaneous emission in these kinds of cavities would also be a very interesting application. Even with current CAMFR capabilities it should already be possible to model spontaneous emission inside oxide-apertured microcavities, but this topic remained largely untouched so far due to time constraints. However, in literature this is considered to be an interesting and challenging problem [Deppe1999]. Also, we could try and develop a more sophisticated model of spontaneous emission in VCSELs near threshold, by going from a single dipole emitter to an ensemble of dipoles, with a spatial distribution described by a carrier diffusion model.

There are also some plans to investigate the use of eigenmode expansion to describe more time-domain like phenomena, e.g. by applying a harmonic balance technique to model non-linear effects.

All of this could increase the flexibility and the possible application domains of CAMFR even further, leading to a very useful set of tools to improve the design of current optical components and to give a numerical proof-of-principle for entirely new and unforeseen devices.

Bibliography

[Abram1998] I. Abram, I. Robert, R. Kuszelewicz, "Spontaneous emission control in semiconductor microcavities with metallic or Bragg mirrors", *IEEE J. Quantum Electron.*, vol. 34, no. 1, pp. 71–98, Jan. 1998.

[Anemogiannis1994] E. Anemogiannis, E. N. Glytsis, T. K. Gaylord, "Efficient solution of eigenvalue equations of optical waveguide structures", *IEEE J. Lightwave Technol.*, vol. 12, no. 12, pp. 2080–2084, Dec. 1994.

[Benisty1998] H. Benisty, R. Stanley, M. Mayer, "Method of source terms for dipole emission modification in modes of arbitrary planar structures", *J. Opt. Soc. Am. A*, vol. 15, no. 5, pp. 1192–1201, May 1998.

[Benisty1999] H. Benisty, C. Weisbuch, D. Labilloy, M. Rattier, C. Smith, T. Krauss, R. De La Rue, R. Houdrée, U. Oesterle, C. Jouanin, D. Cassagne, "Optical and confinement properties of two-dimensional photonic crystals", *IEEE J. Lightwave Technol.*, vol. 17, no. 11, pp. 2063–2077, Nov. 1999.

[Bérenger1994] J.-P. Bérenger, "A perfectly matched layer for the absorption of electromagnetic waves", *Journ. Comput. Phys.*, vol. 114, no. 1, pp. 185–200, 1994.

[Bienstman2000] P. Bienstman, R. Baets, "The RC^{2}LED: a novel resonant-cavity LED design using a symmetric resonant cavity in the outcoupling mirror", *IEEE J. Quantum Elec-*

tron., vol. 36, no. 6, pp. 669–673, June 2000.

[Bienstman2001a] P. Bienstman, R. Baets, J. Vukusiç, A. Larsson, M. Noble, M. Brunner, K. Gulden, P. Debernardi, L. Fratta, G. Bava, H. Wenzel, B. Klein, O. Conradi, R. Pregla, S. Riyopoulos, J.-F. P. Seurin, S. Chuang, "Comparison of optical VCSEL models on the simulation of oxide-confined devices", *IEEE J. Quantum Electron.*, 2001, submitted.

[Bienstman2001b] P. Bienstman, K. De Mesel, C. Sys, K. Naessens, R. Baets, P. Melanen, J. Lindfors, M. Saarinen, M. Pessa, M. Brunner, K. Gulden, F. Salomonsson, M. Hammar, M. Grabherr, M. Miller, J. Mähnss, M. Ghisoni, A. Larsson, L. Hobbs, R. De La Rue, "Comparison of Al(Ga)As oxidation rate of samples from the same wafer on different furnaces", *unpublished*, 2001.

[Bienstman2001c] P. Bienstman, R. Baets, "Optical modelling of photonic crystals and VCSELs using eigenmode expansion and perfectly matched layers", *Opt. Quantum Electron.*, 2001, accepted.

[Bienstman2001d] P. Bienstman, H. Derudder, R. Baets, F. Olyslager, D. De Zutter, "Analysis of cylindrical waveguide discontinuities using vectorial eigenmodes and perfectly matched layers", *IEEE Trans. Microwave Theor. Tech.*, vol. 49, no. 2, pp. 349–354, Feb. 2001.

[Bloch1928] F. Bloch, *Z. Physik*, vol. 52, pp. 555–560, 1928.

[Blomme1996] K. Blomme, *Studie van Greense functies en waveletbasisses bij de modellering van planaire circuits met de momentenmethode (in Dutch)*, PhD thesis, Ghent University, 1996.

[Bockstaele1999] R. Bockstaele, C. Sys, J. Blondelle, B. Dhoedt, I. Moerman, P. Van Daele, P. Demeester, R. Baets, "Resonant cavity LED's optimized for coupling to poly-

mer optical fibers", *IEEE Photon. Technol. Lett.*, vol. 11, no. 2, pp. 158–160, 1999.

[Bond1999] A. Bond, P. Dapkus, J. O'Brien, "Design of low-loss single-mode vertical-cavity surface-emitting lasers", *IEEE J. Sel. Top. Quantum Electron.*, vol. 5, no. 3, pp. 574–581, May 1999.

[Brooks1987] F. P. Brooks, "No silver bullet - essence and accidents of software engineering", *Computer Magazine*, Apr. 1987.

[Brorson1990] S. Brorson, H. Yokoyama, E. Ippen, "Spontaneous emission rate alteration in optical waveguide structures", *IEEE J. Quantum Electron.*, vol. 26, no. 9, pp. 1492–1499, Sep. 1990.

[Chen1988] C. Chen, K. A. Zaki, "Resonant frequencies of dielectric resonators containing guided complex modes", *IEEE Trans. Microwave Theor. Tech.*, vol. 36, no. 10, pp. 1455–1457, 1988.

[Cheng1996] C. Cheng, V. ArbertEngels, A. Scherer, E. Yablonovitch, "Nanofabricated three dimensional photonic crystals operating at optical wavelengths", *Physica Scripta*, vol. 68, pp. 17–20, 1996.

[Chew1990] W. C. Chew, *Waves and fields in inhomogeneous media*, Van Nostrand Reinhold, 1990.

[Chew1994] W. C. Chew, W. H. Weedon, "A 3D perfectly matched medium from modified Maxwell's equations with stretched coordinates", *Microwave Opt. Technol. Lett.*, vol. 7, no. 13, pp. 599–604, Sep. 1994.

[Chew1997] W. Chew, J. Jin, E. Michielssen, "Complex coordinate stretching as a generalized absorbing boundary condition", *Microwave Opt. Technol. Lett.*, vol. 15, no. 6, pp. 363–369, 1997.

[Chow2000] E. Chow, S. Lin, S. Johnson, P. Villeneuve, J. Joannopoulos, J. Wendt, G. Vawter, W. Zubrzyckii, H. Hou, A. Alleman, "Three-dimensional control of light

in a two-dimensional photonic crystal slab", *Nature*, vol. 407, no. 6807, pp. 983–986, Oct. 2000.

[Čtyroký1998] J. Čtyroký, S. Helfert, R. Pregla, "Analysis of a deep waveguide Bragg grating", *Opt. Quantum Electron.*, vol. 30, no. 5-6, pp. 343–358, 1998.

[Dallesasse1990] J. Dallesasse, N. Holonyak, A. Sugg, T. Richard, L. El-Zein, "Hydrolyzation oxidation of $Al_xGa_{1-x}As$/AlGaAs quantum well heterostructures and superlattices", *Appl. Phys. Lett.*, vol. 57, no. 26, pp. 2844–2846, 1990.

[De Neve1997a] H. De Neve, *Study and realization of microcavity based LED's*, PhD thesis, Ghent University, 1997.

[De Neve1997b] H. De Neve, J. Blondelle, P. Van Daele, P. Demeester, R. Baets, G. Borghs, "Planar substrate-emitting microcavity light emitting diodes with 20% external quantum efficiency", In: *Spie Photonics West*, 1997.

[De Neve1997c] H. De Neve, J. Blondelle, P. Van Daele, P. Demeester, R. Baets, G. Borghs, "Recycling of guided mode emission in planar microcavity light emitting diodes", *Appl. Phys. Lett.*, vol. 74, no. 7, pp. 926–928, 1997.

[Delves1967] L. Delves, J. Lynnes, "A numerical method for locating the zeros of an analytic function", *Math. Comp.*, vol. 21, pp. 543–560, 1967.

[Demeulenaere1997] B. Demeulenaere, *Theoretische studie van lichtemissie in halfgeleidercaviteiten met complexe geometrie (in Dutch)*, PhD thesis, Ghent University, 1997.

[Demeulenaere1999] B. Demeulenaere, P. Bienstman, B. Dhoedt, R. Baets, "Detailed study of AlAs-oxidized apertures in VCSEL cavities for optimized modal performance", *IEEE J. Quantum Electron.*, vol. 35, no. 3, pp. 358–367, 1999.

[Deppe1999] D. Deppe, L. Graham, D. Huffaker, "Enhanced spontaneous emission using

quantum dots and an apertured microcavity", *IEEE J. Quantum Electron.*, vol. 35, no. 10, pp. 1502–1508, Oct. 1999.

[Derudder1998] H. Derudder, F. Olyslager, D. De Zutter, "Analysis of waveguide discontinuities using perfectly matched layers", *Electron. Lett.*, vol. 34, no. 22, pp. 2138–2140, 1998.

[Derudder2001a] H. Derudder, *Nieuwe toepassing voor perfect aangepaste lagen in planaire circuits en golfgeleiders (in Dutch)*, PhD thesis, Ghent University, 2001.

[Derudder2001b] H. Derudder, F. Olyslager, D. De Zutter, S. Van den Berghe, "Efficient mode-matching analysis of discontinuities in finite planar substrates using perfectly matched layers", *IEEE Trans. Antennas Propagat.*, vol. 49, no. 2, pp. 185–195, Feb. 2001.

[Dhoedt1996] B. Dhoedt, *Theoretical and experimental study of free-space optical interconnections based on diffractive lens arrays*, PhD thesis, Ghent University, 1996.

[Eitel2000] S. Eitel, S. Fancey, H. Gauggel, K. Gulden, W. Bachtold, M. Taghizadeh, "Highly uniform vertical-cavity surface-emitting lasers integrated with microlens arrays", *IEEE Photon. Technol. Lett.*, vol. 12, no. 5, pp. 459–461, 2000.

[Geels1991] R. Geels, S. Corzine, L. Coldren, "InGaAs vertical-cavity surface-emitting lasers", *IEEE J. Quantum Electron.*, vol. 27, no. 6, pp. 1359–1367, 1991.

[Gerdes1992] J. Gerdes, B. Lunitz, R. Pregla, "Analysis of slab waveguide discontinuities including radiation and absorption effects", *Electron. Lett.*, vol. 28, no. 11, pp. 1013–1015, May 1992.

[Giboney1998] K. S. Giboney, L. B. Aronson, B. E. Lemoff, "The ideal light source for datanets", *IEEE Spectrum*, pp. 43–53, Feb. 1998.

[Grabherr1999] M. Grabherr, M. Miller, R. Jager, R. Michalzik, U. Martin, H. Unold, K. Ebeling, "High-power VCSEL's: Single devices and densely packed 2-D-arrays", *IEEE J. Sel. Top. Quantum Electron.*, vol. 5, no. 3, pp. 495–502, 1999.

[Gralak2000] B. Gralak, S. Enoch, G. Tayeb, "Anomalous refractive properties of photonic crystals", *J. Opt. Soc. Am. A*, vol. 17, no. 6, pp. 1012–1020, June 2000.

[Hadley1991] G. R. Hadley, "Transparent boundary condition for beam propagation", *Opt. Lett.*, vol. 16, no. 9, pp. 624–626, May 1991.

[Haes1996] J. Haes, *Study of beam properties of laser diodes and design of integrated beam expansion structures*, PhD thesis, Ghent University, 1996.

[Hegblom1999] E. Hegblom, N. Margalit, A. Fiore, L. Coldren, "High-performance small vertical-cavity lasers: A comparison of measured improvements in optical and current confinement in devices using tapered apertures", *IEEE J. Sel. Top. Quantum Electron.*, vol. 5, no. 3, pp. 553–560, May 1999.

[Herzinger1993] C. Herzinger, C. Lu, T. DeTemple, W. Chew, "The semiconductor waveguide reflectivity problem", *IEEE J. Quantum Electron.*, vol. 29, no. 8, pp. 2273–2281, 1993.

[Huffaker1999] D. Huffaker, D. Deppe, "Intracavity contacts for low-threshold oxide-confined vertical-cavity surface-emitting lasers", *IEEE Photon. Technol. Lett.*, vol. 11, no. 8, pp. 934–936, 1999.

[Ikegami1972] T. Ikegami, "Reflectivity of mode at facet and oscillation mode in double-heterostructure injection lasers", *IEEE J. Quantum Electron.*, vol. 8, no. 6, pp. 470–476, 1972.

[Itoh1989] T. Itoh, Ed., *Numerical techniques for microwave and millimeter wave passive structures*, John Wiley and sons, 1989.

[Joannopoulos1995] J. D. Joannopoulos, R. D. Meade, J. N. Winn, *Photonic crystals: molding the flow of light*, Princeton University Press, 1995.

[Johnson1999] S. Johnson, S. Fan, P. Villeneuve, J. Joannopoulos, L. Kolodziejski, "Guided modes in photonic crystal slabs", *Phys. Rev. B*, vol. 60, no. 8, pp. 5751–5758, Aug. 1999.

[Klein1998] B. Klein, L. Register, K. Hess, D. Deppe, Q. Deng, "Self-consistent Green's function approach to the analysis of dielectrically apertured vertical-cavity surface-emitting lasers", *Appl. Phys. Lett.*, vol. 73, no. 23, pp. 3324–3326, 1998.

[Labilloy1998] D. Labilloy, H. Benisty, C. Weisbuch, T. Krauss, C. Smith, R. Houdré, U. Oesterle, "High-finesse disk microcavity based on a circular Bragg reflector", *Appl. Phys. Lett.*, vol. 73, no. 10, pp. 1314–1316, Sep. 1998.

[Larson1998] M. Larson, M. Kondow, T. Kitatani, K. Nakahara, K. Tamura, H. Inoue, K. Uomi, "GaInNAs-GaAs long-wavelength vertical-cavity surface-emitting laser diodes", *IEEE Photon. Technol. Lett.*, vol. 10, no. 2, pp. 188–190, Feb. 1998.

[Lazzi1997] G. Lazzi, O. P. Ghandi, "On the optimal design of the PML absorbing boundary conditions for he FDTD code", *IEEE Trans. Antennas Propagat.*, vol. 45, no. 5, pp. 914–916, May 1997.

[Lear1994] K. Lear, R. Sneider, K. Choquette, S. Kilcoyne, J. Figiel, J. Zolper, "Vertical cavity surface emitting lasers with 21% efficiency by metalorganic vapor phase epitaxy", *IEEE Photon. Technol. Lett.*, vol. 6, no. 9, pp. 1053–1055, 1994.

[Lee1986] D. L. Lee, *Electromagnetic principles of integrated optics*, John Wiley and sons, 1986.

[Lee1995] S.-L. Lee, Y. Chung, L. A. Coldren, N. Dagli, "On leaky mode approximations

for modal expansion in multilayer open waveguides", *IEEE J. Quantum Electron.*, vol. 31, no. 10, pp. 1790–1802, 1995.

[Leontovich1948] M. Leontovich, *Investigations of radiowave propagation, part II*, Printing house of the Russian Academy of Sciences, Moscow, 1948.

[Li1996] L. Li, "Formulation and comparison of two recursive matrix algorithms for modeling layered diffraction gratings", *J. Opt. Soc. Am. A*, vol. 13, no. 5, pp. 1024–1035, 1996.

[Loehr1998] J. P. Loehr, *Physics of strained quantum well lasers*, Kluwer academic publishers, 1998.

[Lukosz1977] W. Lukosz, R. Kunz, "Light emission by magnetic and electric dipoles close to a plane interface. I. Total radiated power", *J. Opt. Soc. Am. A*, vol. 67, no. 12, pp. 1607–1619, Dec. 1977.

[Maximov2000] M. Maximov, N. Ledentsov, V. Ustinov, Z. Alferov, D. Bimberg, "GaAs-based 1.3 micron InGaAs quantum dot lasers: A status report", *Journal Electron. Materials*, vol. 29, no. 5, pp. 476–486, May 2000.

[McRae1968] E. McRae, *Surface Science*, vol. 11, p. 479, 1968.

[Meade1993] R. Meade, A. Rappe, K. Brommer, J. Joannopoulos, O. Alerhand, "Accurate theoretical analysis of photonic band-gap materials", *Phys. Rev. B*, vol. 48, no. 11, pp. 8434–8437, Sep. 1993.

[Mekis1996] A. Mekis, J. Chen, I. Kurland, S. Fan, P. R. Villeneuve, J. Joannopoulos, "High transmission through sharp bends in photonic crystal waveguides", *Phys. Rev. Lett.*, vol. 77, no. 18, pp. 3787–3790, Oct. 1996.

[Muller1956] D. Muller, "A method for solving algebraic equations using an automatic computer", *Math. Tables, Aids of Comput.*, vol. 10, pp. 208–215, 1956.

[Nicorovici2001] N. A. Nicorovici, A. A. Asatryan, L. C. Botten, K. Busch, R. C. McPhedran, C. M. de Sterke, P. A. Robinson, G. H. Smith, D. R. McKenzie, A. R. Parker, "Multipole methods for photonic crystal calculations", In: *Proceedings NATO-ASI conference on photonic crystals and light localization*, 2001, submitted.

[Noble1998] M. Noble, J. Loehr, J. Lott, "Quasi-exact optical analysis of oxide apertured microcavity VCSEL's using vector finite elements", *IEEE J. Quantum Electron.*, vol. 34, no. 12, pp. 2327–2339, 1998.

[Noble2000] M. Noble, P. Sotirelis, J. Lott, J. Loehr, "Optical physics of VCSEL tapered oxide apertures", In: R. Binder, P. Blood, M. Osinski, Eds., *Proc. SPIE vol. 3944*, pp. 252–263, 2000.

[Ohiso1998] Y. Ohiso, C. Amano, Y. Itoh, H. Takenouchi, T. Kurokawa, "Long-wavelength vertical-cavity lasers with InGaAsP/InP-GaAs/AlAs DBR's by wafer fusion", *IEEE J. Quantum Electron.*, vol. 34, no. 10, pp. 1904–1913, Oct. 1998.

[Oliner1981] A. A. Oliner, S.-T. Peng, T.-I. Hsu, A. Sanchez, "Guidance and leakage properties of a class of open dielectric waveguides. II. New physical effects", *IEEE Trans. Microwave Theor. Tech.*, vol. 29, no. 9, pp. 855–869, 1981.

[Omar1987] A. S. Omar, K. F. Schünenmann, "Complex and backward-wave modes in inhomogeneously and anistropically filled waveguides", *IEEE Trans. Microwave Theor. Tech.*, vol. 35, no. 3, pp. 268–275, 1987.

[Onischenko2000] A. Onischenko, T. Sale, E. O'Reilly, A. Adams, S. Pinches, J. Frost, J. Woodhead, "Progress in the design and development of AlGaInP visible VCSELs", *IEE proceedings - Optoelectronics*, vol. 147, no. 1, Feb. 2000.

[Peng1981] S. Peng, A. Oliner, "Guidance and leakage properties of open dielectric waveg-

uides. I. Mathematical formulations", *IEEE Trans. Microwave Theor. Tech.*, vol. 29, no. 9, pp. 843–855, 1981.

[Pregla1996] R. Pregla, "The method of lines as generalized transmission line technique for the analysis of multilayered structures", *AEÜ*, vol. 50, no. 5, pp. 293–300, 1996.

[Press1993] W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery, *Numerical recipes in C : the art of scientific computing by William H. Press, Saul A. Teukolsky, William T. Vetterling, Brian P. Flannery*, Cambridge University Press, 1993.

[Rogier2001] H. Rogier, D. De Zutter, "Bérenger and leaky modes in microstrip substrates terminated by a perfectly matched layer", *IEEE Trans. Microwave Theor. Tech.*, 2001, to be published.

[Saarinen2000] M. Saarinen, M. Toivonen, N. Xiang, V. Vilokkinen, M. Pessa, "Room-temperature CW operation of red vertical-cavity surface-emitting lasers grown by solid-source molecular beam epitaxy", *Electron. Lett.*, vol. 36, no. 14, pp. 1210–1211, July 2000.

[Sacks1995] Z. Sacks, D. Kingsland, R. Lee, J. Lee, "A perfectly matched anisotropic absorber for use as an absorbing boundary condition", *IEEE Trans. Antennas Propagat.*, vol. 43, no. 12, pp. 1460–1463, Dec. 1995.

[Sagan1989] H. Sagan, *Boundary and eigenvalue problems in mathematical physics*, Dover Publications Inc., New York, 1989.

[Sale1995] T. E. Sale, *Vertical cavity surface emitting lasers*, Research Studies Press - John Wiley and sons, 1995.

[Sammut1976] R. Sammut, A. Snyder, "Leaky modes on a dielectric waveguide: orthogonalty and excitation", *Appl. Opt.*, vol. 15, no. 4, pp. 1040–1044, 1976.

[Scarmozzino2000] R. Scarmozzino, A. Gopinath, R. Pregla, S. Helfert, "Numerical techniques for

modeling guided-wave photonic devices", *IEEE J. Sel. Top. Quantum Electron.*, vol. 6, no. 1, pp. 150–162, Jan. 2000.

[Schlereth1990] K. Schlereth, M. Tacke, "The complex propagation constant of multilayer waveguides - an algorithm for a personal computer", *IEEE J. Quantum Electron.*, vol. 26, no. 4, pp. 627–630, 1990.

[Schubert1994] E. Schubert, N. Hunt, M. Micovic, R. Malik, D. Sivco, A. Cho, G. Zydzik, "Highly efficient light-emitting diodes with microcavities", *Science*, vol. 265, pp. 943–945, Aug. 1994.

[Seurin2000] J.-F. P. Seurin, G. Liu, S. Chuang, L. Chirovsky, W. Hobson, J. Lopata, "Modal competition in implant-apertured index-guided vertical-cavity surface-emitting lasers", *Appl. Phys. Lett.*, vol. 77, no. 23, pp. 3686–3688, Dec. 2000.

[Smith1991] R. Smith, S. Houde-Walter, G. Forbes, "Numerical determination of planar waveguide modes using the analyticity of the disperision relation", *Opt. Lett.*, vol. 16, no. 17, pp. 1316–1318, 1991.

[Smith1995] R. Smith, S. Houde-Walter, "Leaky guiding in nontransparent waveguides", *J. Opt. Soc. Am. A*, vol. 12, no. 4, pp. 715–724, Apr. 1995.

[Snyder1983] A. Snyder, J. Love, *Optical waveguide theory*, Chapman and Hall, London, 1983.

[Stroustrup1997] B. Stroustrup, *The C++ programming language - third edition*, Addison Wesley, 1997.

[Sudbø1993] A. Sudbø, "Numerically stable formulation of the transverse resonance method for vector mode-field calculations in dielectric waveguides", *IEEE Photon. Technol. Lett.*, vol. 5, no. 3, pp. 342–344, 1993.

[Sztefka1993] G. Sztefka, H. Nolting, "Bidirectional eigenmode propagation for large refractive index steps", *IEEE Photon. Technol. Lett.*, vol. 5, no. 5, pp. 554–557, 1993.

[Taflove1995] A. Taflove, *Computational electrodynamics, the finite-difference time-domain method*, Artech House, 685 Canton Street, Norwood, MA 02062, 1995.

[Teixeira1997] F. Teixeira, W. Chew, "Systematic derivation of anisotropic PML absorbing media in cylindrical and spherical coordinates", *Microwave Guided Wave Lett.*, vol. 7, no. 11, pp. 371–373, 1997.

[van Rossum2000] G. van Rossum, "The Python language website", www.python.org, 2000.

[Vassallo1985] C. Vassallo, *Théorie des guides d'ondes électromagnétiques*, Eyrolles, 61 boulevard Saint-Germain, 75005 Paris, 1985.

[Vassallo1997] C. Vassallo, "1993-1995 Optical mode solvers", *Opt. Quantum Electron.*, vol. 29, no. 2, pp. 95–114, Feb. 1997.

[Wierer1999] J. Wierer, D. Kellogg, J. N. Holonyak, "Tunnel contact junction native-oxide aperture and mirror vertical-cavity surface-emitting lasers and resonant-cavity light-emitting diodes", *Appl. Phys. Lett.*, vol. 74, no. 7, pp. 926–928, Feb. 1999.

[Willems1995] J. Willems, *Modeling and design of integrated optical filters and lasers with a broad wavelength tuning range*, PhD thesis, Ghent University, 1995.

[Xu2000] Y. Xu, R. K. Lee, A. Yariv, "Propagation and second harmonic generation of electromagnetic waves in a coupled-resonator optical waveguide", *J. Opt. Soc. Am. B*, vol. 17, no. 3, pp. 387–399, March 2000.

[Yablonovitch1987] E. Yablonovitch, "Inhibited spontaneous emission in solid-state physics and electronics", *Phys. Rev. Lett.*, vol. 58, no. 20, pp. 2059–2062, May 1987.

[Yablonovitch1991] E. Yablonovitch, "Photonic band structure: the face-centered-cubic case employing nonspherical atoms", *Phys. Rev. Lett.*, vol. 67, no. 17, pp. 2295–2298, Oct. 1991.

[Yariv1999] A. Yariv, Y. Xu, R. Lee, A. Scherer, "Coupled-resonator optical waveguide: a proposal and analysis", *Opt. Lett.*, vol. 24, no. 11, pp. 711–713, June 1999.

[Yeh1977] P. Yeh, A. Yariv, C.-S. Hong, "Electromagnetic propagation in periodic stratified media. I. General theory", *J. Opt. Soc. Am. A*, vol. 67, no. 4, pp. 432–438, Apr. 1977.

[Zaki1988] K. A. Zaki, S. Chen, C. Chen, "Modeling discontinuities in dielectric-loaded waveguides", *IEEE Trans. Microwave Theor. Tech.*, vol. 36, no. 12, pp. 1804–1810, 1988.