FACULTY OF ENGINEERING
AND ARCHITECTURE

Compact and Low-Power Silicon Photonic Transceivers for Inter and Intra Data Center Optical Interconnects

Jochem Verbist

Doctoral dissertation submitted to obtain the academic degree of
Doctor of Electrical Engineering

Supervisors

Prof. Johan Bauwelinck, PhD - Prof. Günther Roelkens, PhD

Department of Information Technology
Faculty of Engineering and Architecture, Ghent University

July 2021

GHENT
UNIVERSITY

FACULTY OF ENGINEERING
AND ARCHITECTURE

Compact and Low-Power Silicon Photonic Transceivers for Inter and Intra Data Center Optical Interconnects

Jochem Verbist

Doctoral dissertation submitted to obtain the academic degree of
Doctor of Electrical Engineering

Supervisors
Prof. Johan Bauwelinck, PhD - Prof. Günther Roelkens, PhD

Department of Information Technology
Faculty of Engineering and Architecture, Ghent University

July 2021

GHENT
UNIVERSITY

ISBN 978-94-6355-498-5
NUR 959
Wettelijk depot: D/2021/10.500/46

Members of the Examination Board

Chair

Prof. Filip De Turck, PhD, Ghent University

Other members entitled to vote

Paraskevas Bakopoulos, PhD, Nvidia, Greece
David Mackey, PhD, mBryonics, Ireland
Marianna Pantouvaki, PhD, imec
Prof. Geert Van Steenberge, PhD, Ghent University
Prof. Xin Yin, PhD, Ghent University

Supervisors

Prof. Johan Bauwelinck, PhD, Ghent University
Prof. Günther Roelkens, PhD, Ghent University

voor papa

Dankwoord

Nu ik er eindelijk in geslaagd ben een tastbare bewijs te produceren (dat groter is dan enkele mm^2) van mijn academische activiteiten de afgelopen jaren, is dit de perfecte gelegenheid om een aantal mensen te bedanken die mij hierin hebben bijgestaan.

Allereerst wil ik mijn promotoren, Johan en Gunther, bedanken. Johan om mij de kans en de ruimte te geven om te doctoreren, als ook de vrijheid om zelf te kunnen kiezen in welke richting ik het onderzoek wou sturen. Gunther om mij te introduceren in de wondere wereld van de fotonica en om de tijd te maken voor korte crash courses in Silicon Photonics en IIIV materialen als ik weer met een licht vertwijfelde blik bij hem op kantoor zat om te bespreken of mijn voorstel überhaupt wel mogelijk was. Verder wil ik ook Guy en Scott bedanken voor de boeiende discussie on- en off-topic doorheen de jaren. Ik heb over veel zaken een duidelijke mening en zij ook, en de boeiendste discussie ontstonden als onze meningen divergeerden. Ik ben ook Jean en Mike dank verschuldigd voor de vlotte ondersteuning op technisch en administratief gebied, als ik weer last-minute met een verzoek of vraag afkwam.

Voorts zou ik graag alle leden van de Photonics Research Group willen bedanken om dit koekoeksjong in jullie stimulerende groep op te nemen. Het grote voordeel van in twee onderzoeksgroepen te zitten (naast al de dubbele administratie, practica begeleiding, group meetings, ...) is dat je veel nieuwe mensen ontmoet. Iets waar ingenieurs, sociale dieren als we zijn, enorm naar uitkijken. Mensen die je kan lastig vallen als je twijfelt of je fotonische vraag zo basaal is dat je vreest dat Gunther met zijn ogen zou rollen als je het rechtstreeks aan hem zou gaan vragen. Mensen waartegen je zonder verpinken kan zeggen: leg het mij uit alsof ik een kleuter ben, en die -quid pro quo gewijs- vervolgens een gelijkaardige vraag aan jouw

stellen over high speed electronics over iets dat zij nooit durfde te vragen. Mensen waar ik enorm veel van heb bijgeleerd en die bovendien zeer aangenaam waren om gewoon met te keuvelen zoals Amin, Jing en Kasper. Sorry for the Dutch, Amin and Jing, I was just saying that I learned a lot from you guys. Most memorably: that I should look down on VCSELs and I should embrace IIIV-on-SOI DFB lasers! En uiteraard Kasper, de valse trage van de PRG, de man met de broze knieën, de bodemloze interesse en de enorme catalogus aan fysische en fotonische kennis die hij bovendien ook nog eens duidelijk weet uit te leggen.

Het grootste deel van mijn tijd heb ik echter besteed met mijn collega's van design, sorry, INTECdesign, sorry, IBCN+, sorry, IDLab! Van de oude garde die nog over "de kooi" sprak in plaats van de meetruimte en op zijn minst nog één experiment gedaan hebben waar het woord duobinary in voor kwam, tot aan het aanstormend jong geweld dat woorden als quantum key distribution rondslingert alsof het niets is. De sfeer was reeds erg goed, maar ik zou durven stellen dat ze met rasse schreden is vooruitgegaan sinds we, na een tot de vroege uurtjes durend kerstfeestje, op geregelde tijdstippen onze frustraties op elkaar konden botvieren met een spelletje Call of Duty (herdoopt tot 'pew pew' door dr. Li).

Daarom mijn dank aan al mijn collega's en/of strijdmakkers: Timothy, Arno, Ramses en Renato (de tofste bazen die iemand zich kan wensen en die mij *helemaal niet* verplicht hebben dit over hun te schrijven), Wouter (aka albinoid aka mijn eeuwige competitor voor employee of the month), Jochen (hij die mij onbedoeld een bijnaam opleverde), Bart (aka wieishandsome?), Koen (Koen Verheyen), Elber, Gert-Jan, Haolin (aka André), Joris (aka JVK aka the special one aka Joris die Van Campenhout), Chia-Yi (chia-knees still make me giggle to this day), Laurens (aka den Bogaert), Laurens (aka José aka Hoe zit het nu weer met die nieuwe PhD regels?), Michiel (aka Michielus aka Kunt ge nog eens even de FPGA komen instellen?), Nish (aka sup-sup) en Cédric (I love uniqorns). Moesten jullie nog een teamactiviteit zoeken: ik heb gehoord dat de meetruimte uitermate geschikt is voor een dartstoernooi.

Tot slot wil ik mijn closest collaborators extra bedanken om mee koffie te gaan drinken als het bij mij even niet meer ging, of omgekeerd, om mij te komen halen om koffie te drinken als het bij hun even niet meer ging. Ik weet oprecht niet of ik doorgezet zou hebben na de eerste twee jaren toen de resultaten wat uitbleven, moest de sfeer onder ons niet zo goed geweest zijn. Ik heb het uiteraard over the fab four: Michael

(aka handsome aka scherp-op-vrijdag), Marijn (aka Marino aka de lach), Joris (aka Hubert aka alle-Saabjes-op-het-droge) en Hannes (aka poepie aka de menselijke metronoom). Bedankt gasten, het was een geweldige rit. Dankzij jullie keerde ik zelfs na de grootste baaldagen met plezier terug naar het labo. Ik kijk nu al met heimwee terug naar onze tijd in de vuurtoren van Zwijnaarde. Laten we echter afspreken dat bij een volgend vrijgezellenweekend als de getuige (onvermijdelijkerwijs) opnieuw in het ziekenhuis belandt, we ditmaal op zijn minst de sleutel van het huisje vragen voor hij afgevoerd wordt.

Ook al leek het niet altijd zo, soms was ik ook buiten het labo te vinden en dan bij voorkeur bij Amaryllis. Ergens ga ik het jammer vinden dat je nu niet meer ga moeten proberen uitleggen aan familie en vrienden wat ik nu eigenlijk doe in Gent, en om jou vervolgens een trots en een beetje twijfelend verhaal te horen brengen over “chips met licht” en “het internet sneller maken”. Ik was niet altijd de gemakkelijkste om mee samen te leven, zekere wanneer ik weer eens ’s avonds wou blijven verder meten omdat ik eindelijk (bijna) alle nodige meetapparatuur voor een paar nachten ter beschikking had. Bovendien bleek de keuze, om na mijn overstap naar BiFast mijn PhD af te werken na de uren, duidelijk niet de makkelijkste weg. Zeker toen er een jaar geleden een klein patatje, Remus genaamd, bij kwam. Bedankt voor al je steun en bedankt voor al onze jaren samen. Een engelengeduld en een portie selectieve doofheid voor mijn uiterst vindingrijke woordspelingen waren nog nooit zo mooi verpakt.

En als laatste bedank ik graag u, de belastingsbetaler, die mij de kans gaf om mijn licht te laten schijnen in nieuwe wereld in de hoop dat ik daar een kleine maar betekenisvolle bijdrage in kon leveren. Ik hoop dat het gelukt is.

Scientia vincere tenebras

Jochem Verbist
(aka Manolo)
Berchem, juni 2021

Table of Contents

List of Figures

Nederlandstalige Samenvatting –Dutch Summary–

Weinig uitvindingen hebben de laatste decennia een dermate grote invloed op ons leven gehad als het internet. De hieruit voortvloeiende stroom aan nieuwe applicaties en communicatiemogelijkheden die tot voor een paar jaar nog praktisch onmogelijk leken (zoals bijvoorbeeld kunnen thuiswerken via video conferences tijdens een wereldwijde pandemie, om maar iets te noemen) voeden de vraag naar meer bandbreedte. Hierdoor stijgt het data verbruik zodanig snel dat de bestaande internet infrastructuur continue moet geüpdatet worden om deze toename te blijven ondersteunen.

Al deze data wordt gehuisvest in data centers waar data verzoeken gelinkt worden aan de gehoste content via een netwerk van switches en servers die onderling verbonden zijn via datalijnen. Door de enorme toename aan data zijn deze data centers ook enorm toegenomen in oppervlakte, met het ontstaan van hyperscale data centers tot gevolg. Deze hyperscale data centers zijn gebouwen ter grootte van een of meerder warenhuizen gevuld met rijen servers, en behoren typisch toe tot tech gianten als Amazon, Google of Microsoft. Hierdoor neemt niet enkel de vraag naar meer bandbreedte per verbinding toe, maar ook het aantal verbindingen en het aandeel aan verbindingen dat meer dan enkele tientallen meters moet afleggen. Als gevolg zijn deze interconnecties de laatste decennia geëvolueerd van puur elektronisch over passieve koperen kabels naar lichtsignalen over een glasvezel kabel, aangezien glasvezel een vele malen groter bandbreedte-afstand product kan bieden dan koper.

Een optische link bestaat uit twee zender-ontvanger paren (ook wel *transceivers* genoemd: een samentrekking van *transmitter* en *receiver*) verbonden door een of meerdere glasvezel kabels. Elke transceiver op zich bevat een lichtbron (een laser), een modulator (die de elektrische data

signalen overbrengt op de optische golf), een fotodiode (die de optische signalen terug omzet naar het elektrisch domein) en elektronica om deze optische componenten aan te sturen.

Naargelang het type link worden er verschillende soorten optische componenten gebruikt. Voor de kortere verbindingen (<50-100 m), bijvoorbeeld tussen een server en een switch in hetzelfde rack, worden voornamelijk direct gemoduleerde lasers gebruikt waarbij de laser zowel dienst doet als lichtbron en als modulator. Om de kosten te drukken, hanteren deze links meestal glasvezels met een bredere kern (multimode fiber). Dit beperkt de maximale bandbreedte-afstand product significant maar laat een veel goedkopere alignatie tussen de optische componenten en de glasvezel toe.

Bij de overige, langere links wordt geopteerd om glasvezel met een kleinere kern (single mode fiber) te gebruiken. Bovendien dient de laser hier uitsluitend als lichtbron en wordt data aangebracht via een externe modulator, wat vele malen hogere hogere bandbreedtes en langere glasvezellengtes toelaat. In deze laatste categorie zijn er twee dominante optische platformen: optische circuits op basis van IIIV halfgeleiders (meestal indium phosphide of InP) of op basis van silicium (silicon photonics of SiP).

Fotonische circuits op basis van silicium hebben als voordeel dat ze kunnen gerealiseerd worden met de bestaande mature CMOS infrastructuur in grote oplages op 200 of 300 mm wafers en met hoge yield. Bovendien bestaat ook de mogelijkheid om elektronica en de fotonica in één enkele chip te integreren, wat interessante mogelijkheden biedt om de interconnecties tussen elektronische en optische componenten in een transceiver verder te verbeteren. Momenteel bestaan bijna alle transceivers uit een aparte optische IC en een of meer elektrische IC die de optische aanstuurt (drivers) en/of uitleest (transimpedantie versterkers of TIAs). Hierdoor is de afstand tussen driver en modulator soms te groot om elektrische golffenomenen te vermijden en dient deze als transmissielijn aangestuurd te worden wat extra vermogen kost.

De meest courante geïntegreerde optische modulatoren zijn spanningsafhankelijke juncties, die ofwel de fase van het licht veranderen (fase modulatoren) of een deel van het vermogen absorberen (elektro-absorptie modulatoren of EAMs). Aangezien de optische interconnecties binnen data centers werken op intensiteitsmodulatie en directe detectie, dienen de opgelegde fasedraaiingen eerst nog naar het amplitude domein gecon-

verteerd te worden. Dit gebeurt in een Mach Zehnder modulator (MZM) topologie: het licht wordt opgesplitst in twee paden, ondergaat een fase modulatie en wordt opnieuw bij elkaar gevoegd waardoor er interferentie en bijgevolg amplitude modulatie optreedt. Omdat fasemodulatie in InP of SiP een aanzienlijk zwakker effect is dan elektro-absorptie, zijn deze MZMs vaak meerdere millimeters lang en/of vereisen ze hoge spanningszwaaien (>2 V_{pp}). Dit leidt er toe dat ze als transmissielijn dienen aangestuurd te worden. Vaak is de karakteristieke impedantie van deze transmissielijn echter beduidend lager dan de conventionele 50 Ω in hogesnelheids elektronica, wat hun vermogenverbruik alleen maar doet toenemen. Tot slot nemen deze modulatoren een erg grote chip oppervlak in beslag, wat niet enkel de kost verhoogt maar stilaan ook een bottleneck wordt om hogere bandbreedtes per vierkante millimeter met het oog op verdere integratie te realiseren.

EAMs daarentegen zijn vaak minder dan 100 µm lang waardoor ze eveneens kunnen aangestuurd worden als een kleine, capacitieve last, en geen 50 Ω terminatie vereisen. Dit type modulator is echter minder courant in SiP dan MZMs, onder meer omdat de indirecte bandgap van bulk SiGe hun werkingsgebied beperkt tot C-band (1530-1565 nm) en/of L-band (1565-1625 nm). De conventionele golflengtes in data centers liggen in de O-band (1260-1360 nm) waar de chromatische dispersie minimaal is. Door deze dispersie gaan de golflengtes in de C- en L-band aan een grotere verschillende snelheid reizen wat het bandbreedte-afstand product van de link beperkt en leidt tot een duidelijke penalty bij data snelheden vanaf 50 Gbaud over een 1 km.

Tijdens dit onderzoek hebben we verscheidene opties onderzocht en experimenteel gedemonstreerd om compacte en laagvermogen SiP transceivers te realiseren. Hierbij streefden we naar topologieën die konden aangestuurd worden met binaire signalen aan lage spanningen (<2 V_{pp}), om een zo groot mogelijke CMOS compatibiliteit te garanderen, op basis van SiGe elektro-absorptie modulatoren op een 220 nm SiP platform.

Bij aanvang van dit onderzoek lag nog niet vast wat het geprefereerde modulatieformaat zou worden voor de (toen) komende generatie 400 GbE transceivers die opereerden aan 100 Gb/s/λ, waardoor ook verschillende modulatie formaten zoals non-return-to-zero (NRZ), electrical duobinary (EDB) en 4-puls amplitude modulatie (PAM-4) werden geëvalueerd.

In een eerste experiment demonstreerden we real-time 100 Gb/s trans-

missie van zowel NRZ als EDB op een enkele GeSi EAM in combinatie met een BiCMOS zender en ontvanger chipset. Dit was de eerste gerapporteerde 100 Gbaud transmissie op SiP zonder het gebruik van digitale signaalverwerking. Door de commerciële optische ontvanger te verwisselen met een GeSi EAM die in maximale absorptie stond daar middel van een hoge biasspanning, werd eveneens de eerste 100 Gbaud per lambda chip-naar-chip optische link in silicium gerealiseerd. Bij NRZ was de maximale afstand van de link omwille van de chromatische distorsie beperkt tot 1 km. Door de lagere bandbreedte eisen van EDB kon deze afstand met behulp van dit modulatieformaat uitgebreid worden tot 2 km.

In een volgend fase werd een nieuwe optische modulator topologie uitgedacht op basis van twee binair aangedreven EAMs in een interferometrische structuur. Hierbij werd ook bekeken wat de voor- en nadelen zijn van een optische of elektrische generatie van het vermogenverchil tussen de minst en meest significante bit. Een demonstrator met twee binaire gestuurde EAMs op basis van een elektrische minst significante bit generatie resulteerde in een beduidende beter oogdiagram dan een enkele multilevel EAM. Dit experiment leidde tot de eerste 128 Gb/s PAM-4 zender op silicium fotonica zonder enige vorm van digitale signaalverwerking of digitaal-naar-analoog convertoren, als ook de eerste real-time PAM-4 SiP link boven de 100 Gb/s waarbij de ontvanger een GeSi EAM was in maximale absorptie.

Met het oog op de volgende generatie 0.8 TbE en 1.6 TbE transceivers, keken we vervolgens naar een implementatie van een optische 4-naar-1 multiplexer in het tijdsdomein om 100 Gbaud NRZ en PAM-4 transmissies te realiseren door gebruik te maken van de reeds bestaande 25 Gbaud drivers en een gepulste laserbron. Dit experiment zorgde voor de hoogste gerapporteerde intensiteitsgemoduleerde transmissiesnelheid op SiP zonder digitale signaalverwerking aan 208 Gb/s PAM-4.

In het laatste reguliere hoofdstuk, richtten we de aandacht op coherente transceivers die communiceren door middel van complexe constellaties zoals QSPK en 16-QAM en onderzochten we hoe deze te realiseren met intensiteitsmodulatoren om zo veel mogelijk van de voordelen uit de voorgaande optische structuren (hoge bandbreedte, compacte vorm, laag vermogen) te kunnen overdragen naar dit type transceiver.

We besluiten dit het boek met een samenvatting van de belangrijkste resultaten, een aantal suggesties voor verbeteringen/alternative versies van

de optische DAC topologie, en een korte vooruitblik op de mogelijkheden zijn voor de volgende generatie(s) interconnects.

English Summary

During the last decades few inventions have had a bigger influence on our lives than the Internet. The resulting stream of new applications and means of communication that may have seemed impractical only a few years ago (e.g. the possibility to work seamlessly from home via high quality video conferences during a worldwide pandemic, to name just one), feed the demand of more bandwidth. Which, in it's turn, paves the way for even newer applications with even higher bandwidth requirements. The rapid and continuous increase in data consumption has led to a permanent need to update the existing internet infrastructure in order to keep supporting this growth.

All this data is housed in so-called data centers where data queries are linked to hosted content via a network of switches and servers, connected with high speed data cables. The enormous growth of the data traffic has also led to emergence of hyperscale data centers: large warehouse scale buildings with several hundred meters of server racks owned, typically owned by one of the tech giant such as Amazon, Google or Microsoft. This evolution has not only pushed the bandwidth per interconnect, but also the total number of interconnects as well as the number of interconnects dat have to reach beyond several tens of meters. Consequently, these links have to transition from pure electronic signaling over passive copper cables to optical signaling over fiber, as fiber supports a much higher bandwidth-distance product than copper.

Such an optical link made up out of two transceivers (a contraction of transmitter and receiver) connected by one or more optical fibers. Each transceiver is consists of a light source (more specifically, a laser), a modulator to convert the electrical data signal to the optical domain, a photodetector to convert that optical signal back to the electrical domain, and the electronics to control these optical components.

Depending on the type of interconnect these functionalities are realized

with different optical components. For short reaches (<50-100), e.g. between a server and a switch in the same rack, directly modulated laser are predominantly used. By modulating the current that flows through the laser diode, the laser acts both as a light source and as modulator. These links usually employ multimode fiber (MMF), i.e. fiber with a wider core. Although this limits the maximal bandwidth-distance product of the link to several tens of meters at the current data rates, it also allows for a much less complex and thus cheaper alignment between the MMF fiber and the optical components.

All other (longer) optical links in data centers use a fiber with a smaller core called single mode fiber (SMF). Now, the laser only serves as the light source for the transceiver module and data is applied through an external modulator. These properties allow data to be transmitted at higher bandwidths and over significantly longer fibers spans. In this dissertation we focus on this last category. Here we can distinguish two main platforms for photonic integrated circuits (PICs): IIIV-based materials (usually Indium Phosphide or InP) or Silicon based (called Silicon Photonics or SiP).

Silicon photonic based PICs have the advantage that they can leverage the existing and mature CMOS infrastructure to fabricate in high volume and high yield on 200 or 300 mm wafers. It is even possible to push the integration further and combine silicon photonics and electronics in a single chip. This monolithic integration offers new ways to improve the IOs between the optical and the electrical domain even further. Currently, almost all transceivers consist of (one or more) optical ICs and one or more electrical ICs with drivers and transimpedance amplifiers which can limit the implementation options. For example, if the high speed contacts of driver cannot be placed close enough to those of modulator the connection between the both often has to be treated as a transmission line, which burns additional power due to the required resistive terminations.

The most common integrated modulators on SiP or InP are voltage dependent diodes which either alter the phase of the light travelling in a waveguide through the diode (i.e. phase modulators) or absorb a fraction of the optical power resulting in intensity modulation (electro-absorption modulators). To achieve intensity modulation with phase modulators, a Mach Zehnder modulator topology is often used: light is split in two paths, undergoes an specific phase modulation, and is recombined leading to interference and, hence, amplitude modulation. As phase modulation

in SiP or InP is typically a much weaker effect than electro-absorption, the modulator has to be made several millimeters long (and consequently large capacitive load) and/or require high voltage swings (>2 V_{pp}) to induce a sufficiently large phase shift. This means that in practice these modulators are usually driven as transmission lines. Moreover, in SiP the characteristic impedances of these transmission lines and terminations are often limited to 25-30 Ω rather than the more conventional 50 Ω in high-speed electronics, resulting in an even higher power consumption. The large chip area required for these MZMs not only often dominates the size of the die and therefore its cost, but it also starts to pose a bottleneck to further increase the bandwidth per square millimeter. A parameter that is especially important for on-board optics, where pluggable optical modules are avoided by bringing the optics even closer to the processing units (e.g. a switch IC).

EAMs on the other hand are very compact modulators. With lengths of less than 100 µm they also can be driven as a small capacitive load and require no explicit resistive termination. On SiP platforms they are, however, less common than MZMs because the indirect bandgap of bulk SiGe based EAMs limits there operating wavelength range to C-band (1530-1565 nm) or L-band (1565-1625 nm). SMF transceivers inside data center conventionally use wavelengths in the O-band (1260-1360 nm) as the chromatic dispersion is minimal in this band. Because of this chromatic distortion longer wavelength such as in C- and L-band will travel at increasingly different speeds through the fiber which fundamentally limits its bandwidth-distance product for intensity modulation. For example, clear penalties are observed for transmission over 1 km for symbol rates of 50 Gbaud an higher.

During this research we have investigated, designed and experimentally demonstrated several ways to realize compact and low-power transceivers based on SiGe electro-absorption modulators on a 220 nm SiP platform, striving for topologies which could be driven with low-swing (<2 V_{pp}) binary signals to ensure maximal CMOS compatibility.

At the start of this project, it had not yet been decided what would be the preferred modulation format for the (then) upcoming generation of 400 Gb/s transceivers operating at 100 Gb/s/λ, so several modulation formats such as non-return-to-zero (NRZ), electrical duobinary (EDB), and 4 pulse amplitude modulation (PAM-4) were evaluated to achieve these these lane rates.

In a first experiment, we demonstrated real-time 100 Gb/s transmission with NRZ and EDB on a single GeSi EAM in combination with an in-house developed BiCMOS transmitter and receiver chipset. This was the first reported 100 Gbaud transmission on silicon photonics without the need for any digital signal processing (DSP). By replacing the commercial optical receiver with another GeSi EAM biased at maximal absorption, we realized the first real-time 100 Gbaud per lambda chip-to-chip optical link with a SiP platform. For NRZ modulation the maximal fiber reach was limited to 1 km due to the high chromatic distortion. However, thanks to its lower bandwidth requirements, this reach could be extended beyond 2 km with EDB modulation, covering the longest possible links inside hyperscale data centers.

In a following phase, a novel optical PAM-4 modulator topology was developed that uses two binary driven amplitude modulators such as EAMs in an interferometric structure biased at 90° to ensure an evenly spaced PAM-4 levels. The generation of the least and most significant bit (LSB and MSB) can be done optically by introducing an power difference into the two arms of the interferometer for the same drive voltage of the EAMs, or electrically by introducing different drive voltage for equal input power levels. A demonstrator of this topology based on the electrically induced LSB-MSB difference demonstrated significantly better PAM-4 eye diagrams than a single multilevel driven EAM. With successful transmission and reception of 128 Gb/s PAM-4 without any need for DSP or power-hungry digital-to-analog convertors, this experiment was the first real-time PAM-4 link capable of at least 100 Gb/s in silicon photonics.

Next, we looked ahead at possible implementations for the next generations of Ethernet transceivers at 0.8 TbE and 1.6 TbE which most likely will need to support 100 Gbaud optical lane rates. An optical 4-to-1 serializer was developed and experimentally validated that achieves 104 Gbaud NRZ and PAM-4 by using four EAMs driven at quarter rate (26 Gbaud), which would allow the reuse of the existing generation of drivers, in combination with a pulsed laser source. With 208 Gb/s PAM-4 transmission we showed the highest reported intensity modulated symbol rates for a silicon-based transmitter without any DSP.

In the final regular chapter, we focused our attention to coherent transceivers. These transceivers communicate through complex modulation formats such as QPSK and 16-QAM which require both amplitude and phase modulation. We investigated how to realize this intensity

modulators such as EAMs in order to port the advantages of the previously demonstrated IMDD structures (high bandwidth, compactness, low power) to this class of optical transceivers.

We conclude this book by summarizing the main results, providing a couple of suggestions for improved/alternative implementations of the optical DAC topology, and a brief look into the possibilities for next generation(s) of data center interconnects.

List of Publications

Publications in International Journals

- J. Lambrecht, **J. Verbist**, H. Ramon, M. Vanhoecke, J. Bauwelinck, X. Yin, G. Roelkens, *Low power (1.5pJ/bit) silicon integrated 106 Gb/s PAM-4 optical transmitter*, IEEE Journal of Lightwave Technology, 38(2), pp. 432-438 (2020). [invited]

- **J. Verbist**, M. Vanhoecke, M. Lillieholm, A. Srinivasan, P. De Heyn, J. Van Campenhout, M. Galili, L. Oxenlowe, X. Yin, J. Bauwelinck, G. Roelkens, *4:1 Silicon Photonic Serialiser for Datacenter Interconnects Demonstrating 104 Gbaud OOK and PAM4 Transmission*, Journal of Lightwave Technology, 37(5), pp. 1498-1503 (2019). [invited]

- **J. Verbist**, J. Lambrecht, M. Verplaetse, A. Srinivasan, P. De Heyn, T. De Keulenaer, R. Pierco, A. Vyncke, J. Van Campenhout, X. Yin, J. Bauwelinck, G. Torfs, G. Roelkens, *Real-Time and DSP-free 128 Gb/s PAM-4 Link Using a Binary Driven Silicon Photonic Transmitter*, Journal of Lightwave Technology, 37(2), pp. 274-280 (2019). [invited]

- **J. Verbist**, J. Lambrecht, M. Verplaetse, J. Van Kerrebrouck, A. Srinivasan, P. De Heyn, T. De Keulenaer, X. Yin, G. Torfs, J. Van Campenhout, G. Roelkens, J. Bauwelinck, *DAC-less and DSP-free 112 Gb/s PAM-4 Transmitter Using Two Parallel Electro- Absorption Modulators*, Journal of Lightwave Technology, 36(5), pp. 1281-1286 (2018). [invited]

- K. Alexander, J. George, **J. Verbist**, K. Neyts, B. Kuyken, D. Van Thourhout, J. Beeckman, *Nanophotonic Pockels Modulators on a Silicon Nitride Platform*, Nature Communications, 9(1), pp. 3444 (2018).

- **J. Verbist**, M. Verplaetse, A. Srinivasan, J. Van Kerrebroeck, P. De

Heyn, P. Absil, T. De Keulenaer, R. Pierco, A. Vyncke, G. Torfs, X. Yin, G. Roelkens, J. Van Campenhout, J. Bauwelinck, *Real-Time 100 Gb/s NRZ and EDB Transmission With a GeSi Electro-Absorption Modulator for Short-Reach Optical Interconnects*, Journal of Lightwave Technology, 36(1), pp. 90-96 (2018). [invited]

- **J. Verbist**, J. Lambrecht, B. Moeneclaey, J. Van Campenhout, X. Yin, J. Bauwelinck, G. Roelkens, *40-Gb/s PAM-4 Transmission Over 40 km Amplifier-less Link Using an O-Band EML and a Sub-5V Silicon Germanium APD* , IEEE Photonics Technology Letters, pp. 2238-2241 (2017).

- **J. Verbist**, J. Zhang, B. Moeneclaey, W. Soenen, J. Van Weerdenburg, R. Van Uden, C. Okonkwo, J. Bauwelinck, G. Roelkens, X. Yin, *A 40 Gbaud QPSK/16-QAM Integrated Silicon Coherent Receiver*, IEEE Photonics Technology Letters, 28(19), pp. 2070-2073 (2016).

- J. Zhang, **J. Verbist**, B Moeneclaey, J. van Weerdenburg, R. Van Uden, H. Chen, J. Van Campenhout, C.M. Okwonko, X. Yin, J. Bauwelinck, G. Roelkens, *Compact Low-Power-Consumption 28-Gbaud QPSK/16-QAM Integrated Silicon Photonic/Electronic Coherent Receiver*, IEEE Photonics Journal, 8(1), pp. 7100110 (2016).

- K. Van Gasse, **J. Verbist**, H. Li, G. Torfs, J. Bauwelinck, G. Roelkens, *Silicon Photonics Radio-Over-Fiber Transmitter Using GeSi EAMs for Frequency Up-Conversion*, IEEE Photonics Technology Letters, 31(2), pp. 181-184 (2019).

- L. Bogaert, J. Van Kerrebrouck, L. Breyne, J. Lambrecht, H. Li, K. Van Gasse, **J. Verbist**, M. Vanhoecke, H. Ramon, S. A. Srinivasan, P. De Heyn, J. Van Campenhout, P. Ossieur, P. Demeester, X. Yin, J. Bauwelinck, G. Torfs, G. Roelkens, *SiGe EAM-Based Transceivers for Datacenter Interconnects and Radio Over Fiber*, in IEEE Journal of Selected Topics in Quantum Electronics, vol. 27, no. 3, pp. 1-13, (2021), doi: 10.1109/JSTQE.2020.3027046.

- K. Van Gasse, L. Bogaert, L. Breyne, J. Van Kerrebrouck, S. Dhoore, C. Op de Beeck, A. Katumba, C.Y. Wu, H. Li, **J. Verbist**, A. Rahim, A. Abbasi, B. Moeneclaey, Z. Wang, H. Chen, J. Van Campenhout, X. Yin, B. Kuyken, G. Morthier, J. Bauwelinck, G. Torfs, G. Roelkens, *Analog Radio-Over-Fiber Transceivers Based on III-V-on-silicon Photonics*, IEEE Photonics Technology Letters, 30(21),

pp. 1818-1821 (2018). [invited]

- K. Van Gasse, J. Van Kerrebrouck, A. Abbasi, **J. Verbist**, G. Torfs, G. Morthier, X. Yin, J. Bauwelinck, G. Roelkens, *III-V-on-Silicon Photonic Transceivers for Radio-over-Fiber links*, Journal of Lightwave Technology, 36(19), pp. 4438-4444 (2018).

- A. Abbasi, **J. Verbist**, L. Abdollahi Shiramin, M. Verplaetse, T. De Keulenaer, R. Vaernewyck, R. Pierco, A. Vyncke, X. Yin, G. Torfs, G. Morthier, J. Bauwelinck, G. Roelkens, *100-Gb/s Electro-Absorptive Duobinary Modulation of an InP-on-Si DFB Laser*, IEEE Photonics Technology Letters, 30(12), pp. 1095-1098 (2018).

- A. Abbasi, L. Abdollahi Shiramin, B. Moeneclaey, **J. Verbist**, X. Yin, J. Bauwelinck, D. Van Thourhout, G. Roelkens, G. Morthier, *III-V-on-Silicon C-Band High-Speed Electro-Absorption Modulated DFB Laser* , Journal of Lightwave Technology, 36(2), pp. 252-257 (2018).

- A. Abbasi, B. Moeneclaey, **J. Verbist**, X. Yin, J. Bauwelinck, G.-H. Duan, G. Roelkens, G. Morthier, *Direct and Electro-Absorption Modulation of a III-V-on-silicon DFB Laser at 56 Gbps*, IEEE Journal of Selected Topics in Quantum Electronics, pp. 1501307 (7 pages) (2017).

- Z. Wang, A. Abbasi, U.D. Dave, A. De Groote, S. Kumari, B. Kunert, C. Merckling, M. Pantouvaki, Y. Shi, B. Tian, K. Van Gasse, **J. Verbist**, R. Wang, W. Xie, J. Zhang, Y. Zhu, J. Bauwelinck, X. Yin, Z. Hens, J. Van Campenhout, B. Kuyken, R. Baets, G. Morthier, D. Van Thourhout, G. Roelkens, *Novel Light Source Integration Approaches for Silicon Photonics*, Laser & Photonics Reviews, 11(4), pp. 1700063 (21 pages) (2017).

- A. Abbasi, S. Keyvaninia, **J. Verbist**, X. Yin, J. Bauwelinck, F. Lelarge, G-H. Duan, G. Roelkens, G. Morthier, *43 Gb/s NRZ-OOK Direct Modulation of a Heterogeneously Integrated InP/Si DFB Laser*, Journal of Lightwave Technology, 35(6), pp. 1235-1240 (2017).

- A. Abbasi, C. Spatharakis, G. Kanakis, N. Sequeira André, H. Louchet, A. Katumba, **J. Verbist**, H. Avramopoulos, P. Bienstman, X. Yin, J. Bauwelinck, G. Roelkens, G. Morthier, *High Speed Direct Modulation of a Heterogeneously Integrated InP/SOI DFB Laser*, IEEE Journal of Lightwave Technolgy, 34(8), pp. 1683-1687 (2016).

- X. Yin, J. Van Kerrebrouck, **J. Verbist**, B. Moenelcaey, X.Z. Qiu, J. Bauwelinck, D. Lanteri, F. Blache, M. Achouche, P. Demeester, *An Asymmetric High Serial Rate TDM-PON with Single Carrier 25 Gb/s upstream and 50 Gb/s downstream*, Journal of Lightwave Technology, 34(2), pp. 819-825 (2015).

- A. Abbasi, **J. Verbist**, J. Van kerrebrouck, F. Lelarge, G. H. Duan, x. Yin, J. Bauwelinck, G. Roelkens, G. Morthier, *28 Gb/s Direct Modulation Heterogeneously Integrated C-band InP/SOI DFB Laser*, Optics Express, 23(20), pp. 26479-26485 (2015).

- G. Roelkens, A. Abbasi, P. Cardile, U.D. Dave, A. De Groote, Y. De Koninck, S. Dhoore, X. Fu, A. Gassenq, N. Hattasan, Q. Huang, S. Kumari, S. Keyvaninia, B. Kuyken, L. Li, P. Mechet, M. Muneeb, D. Sanchez, H. Shao, T. Spuesens, A. Subramanian, S. Uvin, M. Tassaert, K. Van Gasse, **J. Verbist**, R. Wang, Z. Wang, J. Zhang, J. Van Campenhout, Y. Xin, J. Bauwelinck, G. Morthier, R. Baets, D. Van Thourhout, *III-V-on-silicon Photonic Devices for Optical Communication and Sensing*, Photonics, 2(3), pp. 969-1004 (2015).

- H. Chen, **J. Verbist**, P. Verheyen, P. De Heyn, G. Lepage, J. De Cster, P. Absil, B. Moeneclaey, X. Yin, J. Bauwelinck, J. Van Campenhout, G. Roelkens, *25 Gbps 1310 nm Optical Receiver Based on a Sub-5V Waveguide-Coupled Germanium Avalanche Photodiode*, IEEE Photonics Journal, 7(4), pp. 7902909 (2015).

- X. Fu, J. Cheng, Q. Huang, Y. Hu, W. Xie, M. Tassaert, **J. Verbist**, K. Ma, J. Zhang, K. Chen, C. Zhang, Y. Shi, J. Bauwelinck, G. Roelkens, L. Liu, S. He, *5×20 Gbps Heterogeneously Integrated III-V on Silicon Electro-Absorption Modulator Array With Arrayed Waveguide Grating Multiplexer*, Optics Express, 23(14), pp. 18686-18693 (2015).

- H. Chen, **J. Verbist**, P. Verheyen, P. De Heyn, G. Lepage, J. De Coster, P. Absil, X. Yin, J. Bauwelinck, J. Van Campenhout, G. Roelkens, *High Sensitivity 10 Gb/s Si Photonic Receivers Based on a Low-Voltage Waveguide Coupled Ge Avalanche Photodetector*, Optics Express, 23(2), pp. 815-822 (2015).

- M. Verbeke, P. Rombouts, H. Ramon, **J. Verbist**, J. Bauwelinck, X. Yin, G. Torfs, *A 25 Gb/s All-Digital Clock and Data Recovery Circuit for Burst-Mode Applications in PONs*, Journal of Lightwave Technology, vol. 36, no. 8, pp. 1503-1509, IEEE/OSA, (2017).

- H. Ramon, **J. Verbist**, M. Vanhoecke, J. Lambrecht, L. Breyne, G. Torfs, J. Bauwelinck, *A DC-coupled 50 Gb/s 0.064 pJ/bit Thin-Oxide Level Shifter in 28 nm FDSOI CMOS*, IEICE Electronics Express, vol. 15, no. 3, pp. 20171085-20171085, The Institute of Electronics, Information and Communication Engineers, (2018).
- H. Ramon, M. Vanhoecke, **J. Verbist**, W. Soenen, P. De Heyn, Y. Ban, M. Pantouvaki, J. Van Campenhout, P. Ossieur, X. Yin, *Low-Power 56Gb/s NRZ Microring Modulator Driver in 28nm FDSOI CMOS*, IEEE Photonics Technology Letters, vol. 30, no. 5, pp. 467-470, IEEE, (2018).
- J. Lambrecht, H. Ramon, B. Moeneclaey, **J. Verbist**, M. Verplaetse, M. Vanhoecke, P. Ossieur, P. De Heyn, J. Van Campenhout, J. Bauwelinck, *90-Gb/s NRZ Optical Receiver in Silicon Using a Fully Differential Transimpedance Amplifier*, Journal of Lightwave Technology, vol. 37, no. 9, pp. 1964-1973, IEEE/OSA, (2019).
- J. Lambrecht, H. Ramon, B. Moeneclaey, **J. Verbist**, M. Vanhoecke, P. Ossieur, P. De Heyn, J. Van Campenhout, J. Bauwelinck, X. Yin, *A 106-Gb/s PAM-4 Silicon Optical Receiver*, IEEE Photonics Technology Letters, vol. 31, no. 7, pp. 505-508, IEEE, (2019).
- H. Ramon, J. Lambrecht, **J. Verbist**, M. Vanhoecke, S. A. Srinivasan, P. De Heyn, J. Van Campenhout, P. Ossieur, X. Yin, J. Bauwelinck, *70 Gb/s Low-Power DC-coupled NRZ Differential Electro-Absorption Modulator Driver in 55 Nm SiGe BiCMOS*, Journal of Lightwave Technology, vol. 37, no. 5, pp. 1504-1514, IEEE/OSA, (2019).
- H. Li, M. Verplaetse, **J. Verbist**, J. Van Kerrebrouck, L. Breyne, C. Y. Wu, L. Bogaert, B. Moeneclaey, X. Yin, J. Bauwelinck, *Real-Time 100-Gs/s Sigma-Delta Modulator for All-Digital Radio-Over-Fiber Transmission*, Journal of Lightwave Technology, vol. 38, no. 2, pp. 386-393, IEEE/OSA, (2020).

Publications in International Conferences

- **J. Verbist**, M. Lillieholm, J. Van Kerrebrouck, A. Srinivasan, P. De Heyn, J. Van Campenhout, M. Galili, L. Oxenlowe, X. Yin, J. Bauwelinck, G. Roelkens, *104 Gbaud OOK and PAM-4 Transmission Over 1km of SMF Using a Silicon Photonics Transmitter With Quarter-Rate Electronics*, Optical Fiber Communications Conference and Exhibition (OFC 2019), United States, Tu2I.2 (3 pages) (2019). [Top-Scoring]

- **J. Verbist**, J. Lambrecht, M. Verplaetse, S. A. Srinivasan, P. De Heyn, T. De Keulenaer, A. Vyncke, M. Vanhoecke, J. Van Campenhout, X. Yin, G. Roelkens, G. Torfs, *First Real-Time Demonstration of 128Gb/s PAM-4 Transmission Over 1 km SMF Using a Si Photonics Transmitter*, European Conference on Optical Communication, Italy, Mo3I (3 pages) (2018). [Top-Scoring]

- **J. Verbist**, M. Verplaetse, J. Lambrecht, A. Srinivasan, P. De Heyn, T. De Keulenaer, R. Pierco, A. Vyncke, P. Absil, X. Yin, G. Torfs, J. Van Campenhout, G. Roelkens, J. Bauwelinck, *100 Gb/s DAC-less and DSP-free Transmitters Using GeSi EAMs for Short-Reach Optical Interconnects*, Optical Fiber Communication Conference (OFC), United States, pp. W4D.4 (2018). [invited]

- **J. Verbist**, J. Lambrecht, M. Verplaetse, J. Van Kerrebrouck, A. Srinivasan, P. De Heyn, T. De Keulenaer, X. Yin, J. Van Campenhout, G. Roelkens, J. Bauwelinck, *DAC-less and DSP-free PAM 4 Transmitter at 112 Gb/s With Two Parallel GeSi Electro-Absorption Modulators*, European Conference on Optical Communication (ECOC), Sweden, pp. PDP.C.5 (2017). [Post-Deadline Paper]

- **J. Verbist**, M. Verplaetse, A. Srinivasan, P. De Heyn, T. De Keulenaer, R. Pierco, R. Vaernewyck, A. Vyncke, P. Absil, G. Torfs, X. Yin, G. Roelkens, J. Van Campenhout, J. Bauwelinck, *First Real-Time 100-Gb/s NRZ-OOK Transmission Over 2 km With a Silicon Photonic Electro-Absorption Modulator*, Optical Fiber Communication Conference (OFC) - Postdeadline Papers, United States, pp. 1-3, Th5C.4 (2017). [Post-Deadline Paper]

- **J. Verbist**, M. Verplaetse, A. Srinivasan, P. De Heyn, T. De Keulenaer, R. Vaernewyck, R. Pierco, A. Vyncke, P. Verheyen, S. Balakrishnan, G. Lepage, M. Pantouvaki, P. Absil, X.Yin, G. Roelkens, G.

Torfs, J. Van Campenhout, J. Bauwelinck, *Real-Time 100 Gb/s NRZ-OOK Transmission With a Silicon Photonics GeSi Electro-Absorption Modulator*, IEEE Optical Interconnects Conference (OI) 2017, United States, pp. 29-30 (2017).

- **J. Verbist**, J. Zhang, B. Moeneclaey, W. Soenen, J. Van Weerdenburg, R. Van Uden, C. Okonkwo, X. Yin, G. Roelkens, J. Bauwelinck, *A 40 GBaud Integrated Silicon Coherent Receiver* , European Conference on Integrated Optics (ECIO), Poland, pp. ECIO/o-22 (2016). [Best student paper]

- **J. Verbist**, H. Chen, B. Moeneclaey, J. Van Campenhout, X. Yin, J. Bauwelinck, G. Roelkens, *A 1310 Nm Sub-5V Ge APD Based Optical Receiver in a Silicon Platform for 32 Gbps PAM-4 Transmission Over 40 km*, Proceedings Symposium IEEE Photonics Society Benelux, Belgium, pp. 91-94 (2016).

- **J. Verbist**, J. Zhang, B. Moeneclaey, J. van Weerdenburg, R. van Uden, C. Okonkwo, X. Yin, J. Bauwelinck, G. Roelkens, *112 Gbit/s Single-Polarization Silicon Coherent Receiver With Hybrid-Integrated BiCMOS Linear TIA*, 41st European Conference on Optical Communication (ECOC), Spain, P.4.15 (2015).

- J. Lambrecht, **J. Verbist**, H. Ramon, M. Vanhoecke, B. Moeneclaey, L. Bogaert, P. Ossieur, J. Van Campenhout, J. Bauwelinck, G. Roelkens, X. Yin, *53 GBd PAM-4 DAC-less Low-Power (1.5pJ/Bit) Silicon Integrated Transmitter*, European Conference on Optical Communication (ECOC), United States, pp. 4459073 (4 pages) (2019).

- K. Van Gasse, A. Abbasi, M. Shahin, **J. Verbist**, J. Van Kerrebrouck, G. Torfs, B. Moeneclaey, J. Bauwelinck, X. Yin, G. Roelkens, G. Morthier, *Silicon Photonic Radio-Over-Fiber Transceivers and Microwave Photonic Up-Converters*, PhotonIcs & Electromagnetics Research Symposium (PIERS), Italy, pp. 2693 (2019).

- K. Van Gasse, **J. Verbist**, H. Li, G. Torfs, J. Bauwelinck, G. Roelkens, *EAM-based Microwave Mixer Implemented in Silicon Photonics*, European Conference on Integrated Optics (ECIO 2019), Belgium, pp. T.B2.4 (2019).

- L. Breyne, **J. Verbist**, J. Zhang, P. Demeester, G. Roelkens, G. Torfs, *Flexible Integrated Silicon Photonics Mm-Wave Frequency Up Conversion Using GeSi EAMs*, European Conference on Integrated

Optics (ECIO 2019), Belgium, (2019).

- L. Bogaert, **J. Verbist**, K. Van Gasse, G. Torfs, J. Bauwelinck, G. Roelkens, *Germanium Photodetector With Monolithically Integrated Narrowband Matching Network on a Silicon Photonics Platform*, European Conference on Integrated Optics (ECIO 2019), 21, Belgium, pp. W.P01.11 (2019).
- A. Abbasi, **J. Verbist**, L. Abdollahi Shiramin, M. Verplaetse, T. De Keulenaer, R. Pierco, A. Vyncke, G. Torfs, G. Morthier, J. Bauwelinck, G. Roelkens, *100 Gb/s Duobinary Electro-Absorption Modulation of a Heterogeneously Integrated InP-on-Si DFB Laser Diode*, European Conference on Integrated Optics (ECIO, Spain, pp. Th.1.A.3-HRP (2018).
- G. Muliuk, K. Van Gasse, M. Shahin, **J. Verbist**, A. J. Trindade, B. Corbett, D. Van Thourhout, G. Roelkens, *4×25Gbit/s Silicon Photonics Tunable Receiver Using Transfer Printed III-V Photodiodes*, IEEE Photonics Conference (IPC), United States, pp. TuA1.2 (169-170) (2018).
- A. Abbasi, B. Moeneclaey, **J. Verbist**, X. Yin, J. bauwelinck, G. Roelkens, G. Morthier, *56 Gb/s Direct Modulation of an InP-on-Si DFB Laser Diode*, Optical Interconnect Conference (OI), United States, (2017).
- A. Abbasi, H. Chen, **J. Verbist**, X.Yin, Y. Bauwelinck, G. Roelkens, G. Morthier, *Chirp Managed Optical Link Based on a Directly Modulated InP/Si DFB Laser*, European Conference on Integrated Optics (ECIO) 2017, Netherlands, pp. T5.5 (2017).
- A. Abbasi, B. Moeneclaey, **J. Verbist**, X. Yin, J. Bauwelinck, G. Roelkens, G. Morthier, *56 Gb/s Electro-Absorption Modulation of a Heterogeneously Integrated InP-on-Si DFB Laser Diode*, Optical Fiber Communication Conference (OFC), United States, Th4G.2 (3 pages) (2017).
- J. Sarmiento-Merenguel, **J. Verbist**, K. Van Gasse, J. Zhang, B. Moeneclaey, J. Bauwelinck, X. Yin, R. Halir, A. Ortega-Monux, I, Molina-Fernandez, G. Roelkens, *Demonstration of Silicon-On-Insulator Coherent Receiver for Radio-Over-Fiber Applications*, European Conference on Integrated Optics (ECIO), Netherlands, pp. M7.4 (2017).

- A. Abbasi, H. Chen, **J. Verbist**, X. Yin, J. Bauwelinck, G. Roelkens, G. Morthier, *Toward Si Photonics Based Transceivers Using Directly Modulated Heterogeneously Integrated DFB Lasers*, Proceedings Symposium IEEE Photonics Society Benelux, Belgium, pp. 11-14 (2016).
- A. Abbasi, **J. Verbist**, X. Yin, F. Lelarge, G. H. Duan, J. Bauwelinck, G. Roelkens, G. Morthier, *Enhanced Modulation Bandwidth of Heterogeneously Integrated III-V-on-silicon DFB Laser for 40 Gb/s NRZ-OOK Direct Modulation*, International Semiconductor Laser Conference (ISLC), Japan, (2016).
- G. Morthier, A. Abass, **J. Verbist**, S. Keyvaninia, X. Yin, F. Lelarge, G-H. Duan, J. Bauwelinck, G. Roelkens, *High-Speed Directly Modulated Heterogenously Integrated InP/Si DFB Laser*, 42nd European Conference and Exhibition on Optical Communications (ECOC), Germany, pp. 148-150 (2016).
- A. Abbasi, **J. Verbist**, X. Yin, F. Lelarge, G-H. Duan, J. Bauwelinck, G. Roelkens, G. Morthier, *Above 40 Gb/s Direct Modulation of a Heterogeneously Integrated III-V-on-silicon DFB Laser*, Latin America Optics and Photonics Conference (LAOP 2016), Colombia, LYu2D3 (2016).
- G. Morthier, A. Abbasi, M. Shahin, **J. Verbist**, G. Roelkens, *High Speed Modulation of InP Membrane DFB Laser Diodes*, International Conference on Transparent Optical Networks (ICTON), Italy, pp. Mo.C5.3 (2016).
- G. Morthier, A. Abbasi, **J. Verbist**, S. Keyvaninia, X. Yin, F. Lelarge, G.H. Duan, J. Bauwelinck, G. Roelkens, *High-Speed Directly Modulated Heterogeneously Integrated InP/Si DFB Laser*, European Conference on Optical Communication, Germany, pp. 148-150 (2016).
- A. Abbasi, C. Spatharakis, G. Kanakis, N. S. André, H. Louchet, A. Katumba, **J. Verbist**, X. Yin, J. Bauwelinck, H. Avramopoulos, G. Roelkens, G. Morthier, *PAM-4 and Duobinary Direct Modulation of a Hybrid InP/SOI DFB Laser for 40 Gb/s Transmission Over 2 km Single Mode Fiber*, Optical Fiber Communication Conference (OFC), United States, pp. M2C.6 (2016).
- H. Chen, **J. Verbist**, P. Verheyen, P. De Heyn, G. Lepage, J. De Coster, P. Absil, X. Yin, J. Bauwelinck, J. Van Campenhout, G. Roelkens, *Low-Voltage Waveguide Ge APD Based High Sensitivity*

10Gb/s Si Photonic Receiver, (41st European Conference on Optical Communication) ECOC, Spain, Tu.1.3.4 (2015).

- A. Abbasi, **J. Verbist**, J. Van kerrebrouck, F. Lelarge, G. H. Duan, G. Roelkens, G. Morthier, *28 Gb/s Direct Modulation Heterogeneously Integrated InP/Si DFB Laser*, European Conference on Optical Communication (ECOC), Spain, We.2.5.2 (2015).
- X. Fu, Q. Huang, Y. Hu, M. Tassaert, **J. Verbist**, J. Cheng, K. Ma, J. Zhang, K. Chen, C. Zhang, Y. Shimura, J. Bauwelinck, G. Roelkens, L. Liu, S. He, *5×20Gbps III-V on Silicon Electroabsorption Modulator Array Heterogeneously Integrated With a 1.6nm Channel Spacing Silicon AWG*, CLEO: 2015 - Laser Science to Photonic Applications, United States, STu4F.2 (2015).
- X. Yin, **J. Verbist**, T. De Keulenaer, B. Moenelcaey, J. Verbrugghe, X.Z. Qiu, J. Bauwelinck, *25Gb/s 3-Level Burst-Mode Receiver for High Serial Rate TDM-PONs*, Optical Fiber Communications Conference and Exhibition (OFC), Th4H.2, (2015).
- A. Abbasi, L. A. Shiramin, B. Moeneclaey, **J. Verbist**, X. Yin, J. Bauwelinck, D. Van Thourhout, G. Roelkens, G. Morthier, *2× 56 Gbps Electroabsorption Modulated III-V-on-Silicon DFB Laser*, 2017 European Conference on Optical Communication (ECOC), 01-Mar, IEEE, 2017.
- H. T. Chen, **J. Verbist**, P. Verheyen, P. De Heyn, G. Lepage, J. De Coster, P. Absil, B. Moeneclaey, X. Yin, J. Bauwelinck, *Sub-5V Germanium Waveguide Avalanche Photodiode Based 25 Gb/s 1310 Nm Optical Receiver*, Asia Communications and Photonics Conference, AM1B. 4, Optical Society of America, 2015.
- H. Ramon, J. Lambrecht, **J. Verbist**, M. Vanhoecke, S. A. Srinivasan, P. De Heyn, J. Van Campenhout, P. Ossieur, X. Yin, J. Bauwelinck, 70 Gb/s 0.87 pJ/bit GeSi EAM Driver in 55 Nm SiGe BiCMOS, 2018 European Conference on Optical Communication (ECOC), 01-Mar, IEEE, 2018.
- H. Li, M. Verplaetse, **J. Verbist**, J. Van Kerrebrouck, L. Breyne, C. Y. Wu, L. Bogaert, X. Yin, J. Bauwelinck, P. Demeester, *Real-Time 100-Gs/s Sigma-Delta All-Digital Radio-Over-Fiber Transmitter for 22.75-27.5 GHz Band*, Optical Fiber Communication Conference (OFC), Th1F. 4, Optical Society of America, 2019.

- X. Yin, J. Lambrecht, G. Coudyzer, **J. Verbist**, H. Ramon, P. Ossieur, G. Torfs, J. Bauwelinck, *Electronic Circuits for High Speed PON Beyond 25G*, 2019 Optical Fiber Communications Conference and Exhibition (OFC), 01-Mar, IEEE, 2019.
- G. Roelkens, **J. Verbist**, K. Van Gasse, M. Verplaetse, J. Van Kerrebrouck, J. Lambrecht, A. Abbasi, B. Moeneclaey, A. Srinivasan, P. De Heyn, *High Speed Silicon Photonic Transceivers*, The European Optical Society Biennial Meeting 2018 (EOSAM 2018), pp. 1, 2018.
- Y. Ban, **J. Verbist**, M. Vanhoecke, J. Bauwelinck, P. Verheyen, S. Lardenois, M. Pantouvaki, J. Van Campenhout, *Low-Voltage 60Gb/s NRZ and 100Gb/s PAM4 O-Band Silicon Ring Modulator*, 2019 IEEE Optical Interconnects Conference (OI), 01-Feb, Ieee, 2019.
- J. Lambrecht, H. Ramon, B. Moeneclaey, **J. Verbist**, P. Ossieur, P. De Heyn, J. Van Campenhout, J. Bauwelinck, X. Yin, *56-Gb/s Silicon Optical Receiver Using a Low-Noise Fully-Differential Transimpedance Amplifier in SiGe BiCMOS*, 2018 European Conference on Optical Communication (ECOC), 01-Mar, IEEE, 2018.
- J. Van Kerrebrouck, H. Li, J. Verbist, G. Torfs, J. Bauwelinck, and P. Demeester, *Nonlinear MIMO impulse responses determination using pseudo random sequences, and nonlinear compensation*, in Advances in Wireless and Optical Communications (RTUWO), Riga, Latvia, pp. 32–35, 2017.

1 Introduction

1.1 And then there was the Internet

Since its creation in the late 1960s the Internet has drastically changed many aspects of our every day lives, and this trend only seems to accelerate. As with many disruptive technologies, the Internet has grown way beyond its humble beginnings as a network between a couple of university campuses [1] into a giant global network that reaches more than half of the population of the world[2].

After the development of the world-wide web (WWW) in 1989 where users could easily find and access information through webpages [3], the Internet quickly found its way to the general public. Only two years after the first web page went 'on-line', the global internet traffic amounted to 100 GB per day. Ten years later this traffic had already risen spectacularly to 100 GB per second, and is expected to grow to more than 150000 GB/s in 2022 [4]. It is often insightful to concretize these numbers into more familiar and tangible figures (e.g. the equivalent of 31915 DVDs, or 214286 CDs or even 107142858 floppy disks of data per second in 2022). Ironically, the aggressive growth of the internet (Fig. 1.1) will make these types of comparisons feel outdated soon after they are written down as it becomes much more practical to store everything 'in the cloud' rather than on

physical devices in our offices or in our homes.

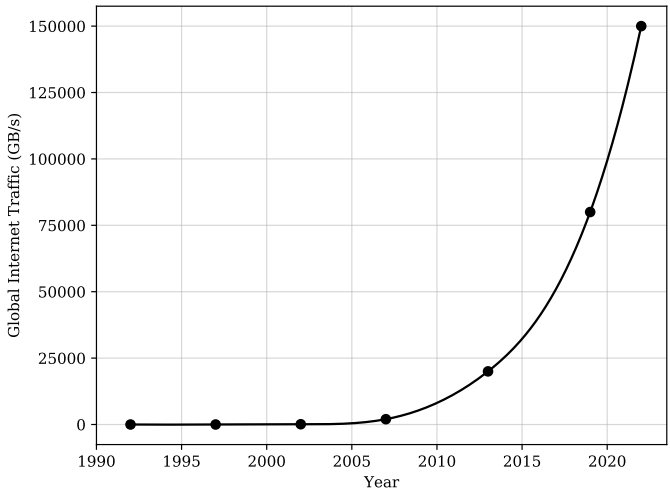

Figure 1.1: Global internet traffic over time (sourced from [4]).

The widespread access to broadband internet has led to the emergence of new applications such as video streaming service (Netflix, YouTube,...), cloud storage (OneDrive from Microsoft, Google Drive, iCloud from Apple,...), and social media (Facebook, Instagram, WhatsApp, Twitter, WeChat,...). Many of these applications have been able to gather an enormous user base as shown in Fig. 1.2, and have become significant contributors to the global internet traffic in only a few years time. Video streaming alone accounted for 60% of the global traffic in 2019 [5] and is expected to grow to 70 or even 80% by 2022[4].

With the majority of the total internet traffic coming from video streaming, it comes as no surprise that the PC is no longer the main consumer of traffic. On the contrary, TVs and smartphones already take up twice the amount of traffic and smartphone users alone will account for half of the internet traffic by 2022 [4]. And with the emergence of the internet of things (IoT), the number of interconnected devices will increase dramatically as well, fueling the ever continuing growth of the internet (an estimated 7% of the total traffic by 2022) even further[4].

As the success of current applications drives the demand for more bandwidth, so will the availability of bandwidth lead to emergence of new applications which in turn will push the demand for more bandwidth, reigniting the cycle once again [7].

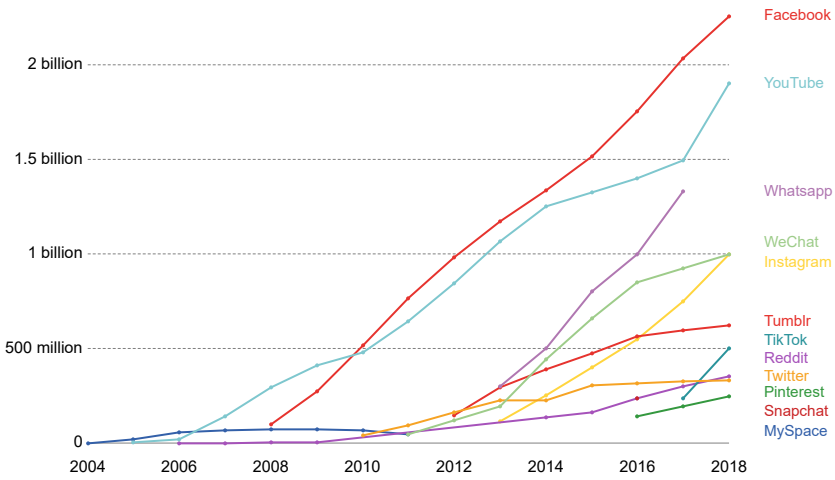

Figure 1.2: Historic overview of the amount of users of the most popular social media platforms starting from with MySpace as the first platform to reach 1 million users in 2004 [6]. Estimates correspond to monthly active users.

1.2 A Network of Networks

This continued growth has put tremendous strain on the existing internet infrastructure. In order to accommodate the request for more and more bandwidth, the underlying networks have had to steadily increase their capabilities.

Figure 1.3 shows the internet architecture as a three-layer hierarchy. In the center, we find the core network that interconnects continents and countries with each other ($\sim$ 1000 km links). Densely populated areas, such as large cities, are connected by the metropolitan (or metro) networks ($\sim$ several 100 km's). At the bottom of the hierarchy, are the access networks ($\leq$ 50 km), which provide the last fiber link before servicing the end-user.

For most people, the relatively short connection after the last piece of fiber and their homes is typically the most familiar one (although fiber-to-the-home deployment is on the rise). For this fixed wired network there are two common options in Belgium: Digital Subscriber Line (DSL) over the 'old' twisted pair telephone lines and the Data Over Cable Service Interface Specification (DOCSIS) over a coaxial cable (often referred to as the 'TV cable'). The wireless access (apart from WiFi) is provide by 4G -and soon

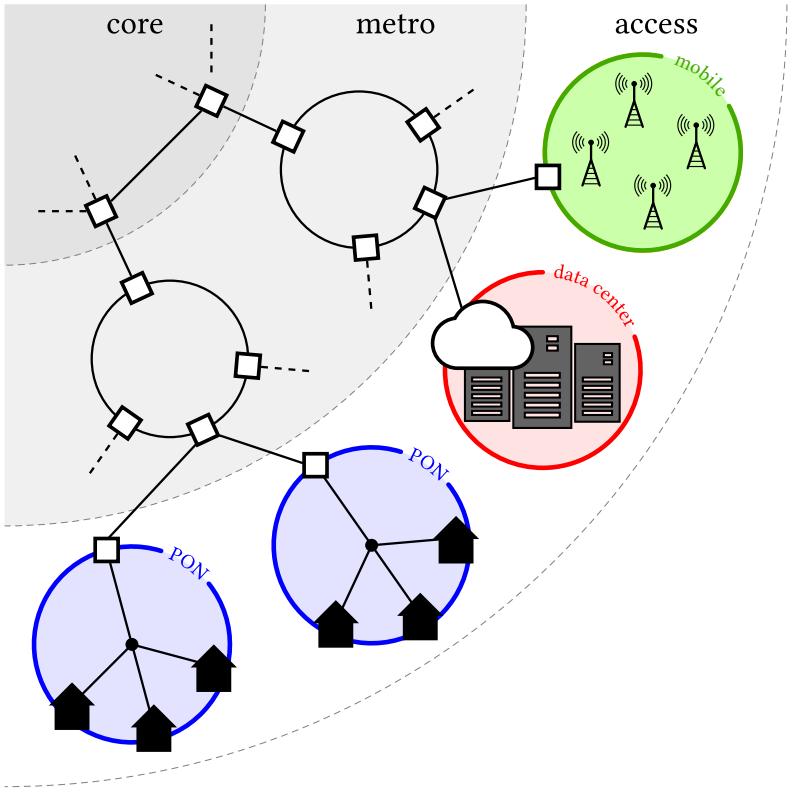

Figure 1.3: The architecture of the internet divided in three layers: core, metro and access.

5G- capable cell towers which are also connect with fiber. Apart from a cell tower or the grey cabins in a couple of streets, the end-user can also be a big company building, university campus or data center (DC).

These boundaries are rather fluid and change depending on the situation. The categories mainly serve to provide a rule-of-thumb framework when discussing the requirements of a link in core, metro or access segment of the internet. Data centers for example can be situated in the access part, but often connect directly to a metro (or even core) hub due to the shear size of their data traffic. As they have become such an integral part of the internet over the last decade, they are now an important category in their own right with a market value of more than 20 billion USD in 2018 (more than e.g. the global music industry [8]), and projected to more than triple in size by 2025 [9].

1.3 Data Centers: Putting the Big in Big Data

Data centers have existed for quite some time as a centralized place with tens or hundreds of servers which provide computational processing power and storage for large companies, universities or stock exchanges.

Until last decade, the telecom industry had been the main driver in the development of optical transceivers, for example with the deployment of passive optical networks (PONs). Thanks to their past efforts most (of the densely populated parts) of the industrialized world had access to reliable broadband internet. Due to high cost associated with placing or replacing parts of PONs, these networks are typically deployed in such a way that would not require major upgrades in the foreseeable future. However, this also meant that there was no immediate need to push the transceivers to much higher data rates, as a reduction of cost (either through cheaper client side transceivers or by increasing the amount of transceivers that could be serviced by one optical line terminal) was probably more important from a telecom operator point of view.

As more and more people found access to broadband internet both at home and on their mobile devices, the aggregate data consumption soared in large part to the boom of many cloud based services. To handle this exponential growth in data, the data centers of the major players in these domains had to rapidly adapt to keep up. This led to the emergence of hyperscale data centers, warehouse sized buildings filled with (predominantly optically) interconnected racks of servers and switches (Fig. 1.4).

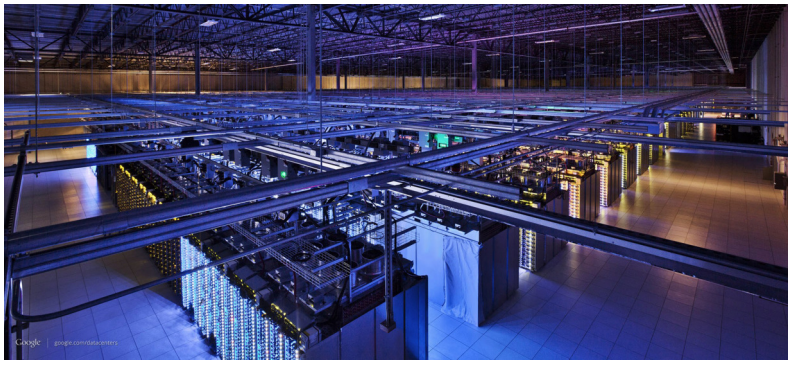

Figure 1.4: Google data center in Hamina, Finland.

Fig. 1.5 shows the internal architecture of a conventional DC, i.e. how all the servers and switches are connected with one another. Although the

most efficient DC architecture is an active research domain on its own, as -for example- not all applications can tolerate the additional latency by traversing up and down the rack-switch tree to send data from one rack to another[10, 11].

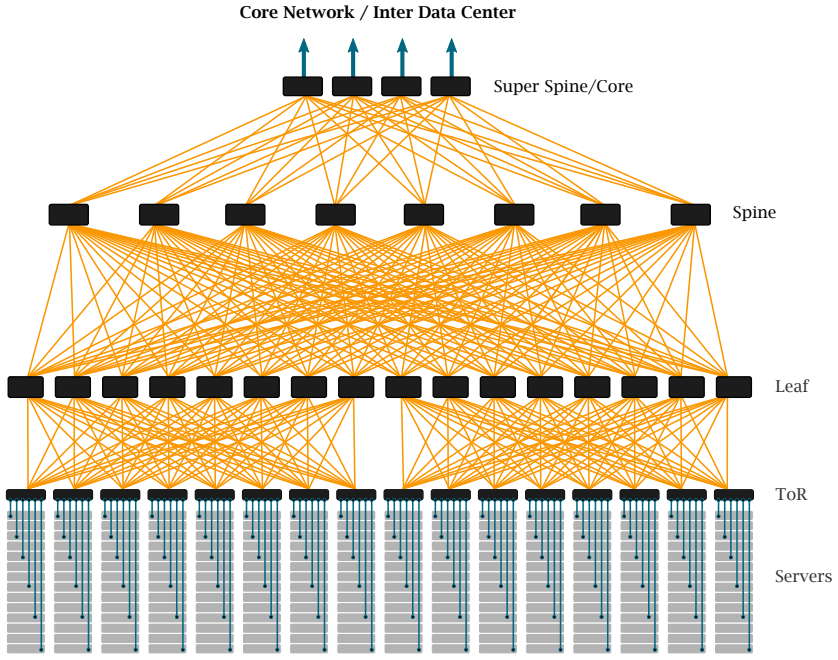

Figure 1.5: An interconnection architecture of switches and racks in a data center (ToR: top of rack switch).

Each switch holds several tens of pluggable modules, all housing one of more transceivers (a contraction of *transmitter* and *receiver*) which readies the data coming from the switch for transmission over the physical channel (i.e. the optical fiber) as well as the reverse operation for data coming from the channel. The components in a typical optical data center link are depicted in Fig. 1.6.

Figure 1.6: A typical optical data center interconnect link.

These hyperscale DCs are typically owned and operated by the company it supports. The biggest players in this field are AWS (Amazon web services), Google, Microsoft, Facebook, Apple, and Baidu. Most of these hyperscale data centers are situated in the USA, but as trend is to build these DCs closer and closer to the end-user, the fraction of DCs in Europe and Asia has grown steadily. In Belgium, for example, Google has a data center in St. Ghislain which can be seen in Fig. 1.7. Since the completion of the first DC in 2010, the site has been extended with two additional DCs and last year Google announced its plans for a fourth DC with an investment of 600 million USD [12].

Figure 1.7: Example of a Hyperscale data center of Google in St. Ghislain, Belgium. Since the completion of the first DC in 2010 in St. Ghislain, Google has extended the site with two new DCs and has announced a fourth DC in June 2019 with an investment of 600 million USD [12].

The success of the internet is perhaps most apparent in the evolution of the five biggest global companies by market capitalization. In 2018, these positions were filled by Apple, Amazon, Alphabet (the parent holding of Google), Microsoft and Facebook, replacing titans such as Shell and ExxonMobil [13]. Data has become the oil of the 21st century.

The number of hyperscale data centers has increased steadily to accommodate the shift of local to in-the-cloud data usage from around 100 in 2010 to more than 500 in 2019[14]. Nevertheless, only increasing the amount of

servers will not suffice to sustain the increased data traffic as for every bit sent from/to a data center 6 bits are sent within or between data centers[4]. As Fig. 1.8 indicates, the traffic between the user and the DC only accounts for 15% of the total DC traffic. More than 85% of the generated bits never leave the DC networks. This means that apart from building more or larger data centers, the data volume per the link inside those buildings will also have to scale drastically.

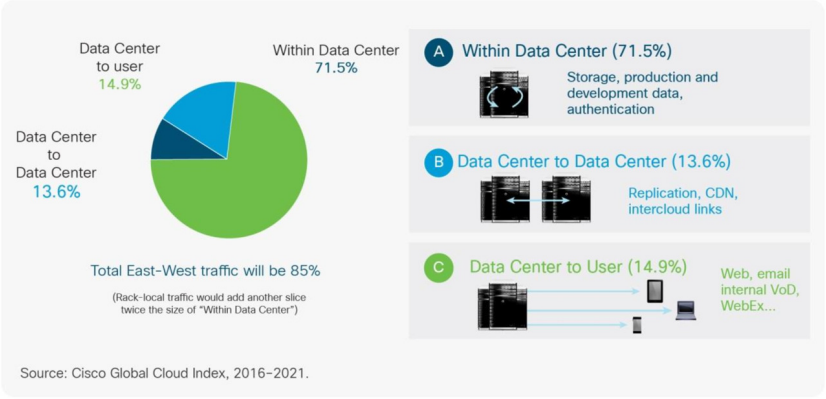

Figure 1.8: Fraction of the data that travels within or between DCs (East-West traffic) and between the DC and the user (North-South traffic)[4].

1.3.1 Intra DCIs

Inside a single data center there are three types of links: copper-based (<5m), multi-mode fiber(<50-100m) and single-mode fiber (>100m). With the increasing data rates copper is being pushed to shorter and shorter distances and will likely disappear in the near future. In general, multi-mode fiber links are much more prevalent than the single-mode ones as there are many more short intra-rack connections. Unsurprisingly, these interconnects are also much cheaper than single-mode based interconnects. However, next to an intrinsically limited link reach, multimode optics also struggle to keep up with the bandwidth requirements for the next generation links. This opens the door for more performant but more costly single-mode optics.

Despite their enormous dimensions, standardized data center interconnects (DCIs) up to 2 km (e.g. the FR links in the Ethernet naming convention [15]) should be able to easily support even the biggest DC building in terms of fiber length. DCIs ranging up to 2 km are typically

referred to as *intra-DCIs* as they are used almost exclusively within a single DC.

1.3.2 Inter DCIs

Apart from connections inside the data center, there are also specific fiber links outside of the data center to connect e.g. two data centers or two buildings within a data center campus with each other. As the end-points in these links can be several tens of kilometers apart, the reach of intra-DCIs is insufficient and a different type of DCI is used, called *inter-DCIs*. As data centers have expanded over the last couple of years, once an implementation has grown beyond a certain limit (be it number of servers, megaWatts or square feet) it looses the economies of scale. Even worse, the consequences of a single DC failure (e.g. through a natural disaster or a power outage) become problematic for the whole system. That is why there is also a trend to create clusters of DCs in specific (often highly populated) geographical areas[16]. Data can be mirrored in several facilities for redundancy without suffering large latency penalties, and yet still distant enough to minimize the change of a joint failure due to an external event. These inter-DCIs typically range between 10 and 80 km as 80 km is deemed the maximal tolerable latency penalty between two data centers [17, 18].

Currently, the inter-DC traffic is already on par with the traffic between the user and the DC (Fig. 1.8) and it is -yet again- expected to grow more rapidly in the coming years. These optical links are especially challenging as they combine the distances from the classic passive optical networks (PONs) in the access segment with the extreme data rates and low-cost profile ($\sim$ 1\$/ Gb/s) of intra-DCIs.

1.3.3 Power consumption

Unfortunately, the power consumption of new high-capacity generation of optical interconnects cannot scale at the same rate as the data traffic both for economical, environmental and practical reasons. Hyperscale DCs often require their own dedicated energy production to minimize the burden on the local power grid [19] as they easily consume anywhere from 15 MW to 150 MW.

In 2018, DCs were responsible for 1% of the global electricity demand (and 0.3% of the global carbon footprint), and some estimate that this figure might rise to 3-13% by 2030 according to a study from 2015 [13]. However, more recent projections from the International Energy Agency

(IEA) paint a more optimistic picture for the DC industry with estimates around 3% [20]. One of the reasons for this outlook from the IEA is that during the past decade, the energy demand for all data centers has not changed drastically, even though the data traffic has (e.g. DC workloads has more than tripled from 2014 to 2020). The main driver for this stable energy consumption in the last five years was the shift from regular data centers to large hyperscale data centers as can be seen in Fig. 1.9. These hyperscale data centers are much more efficient per gigabit as they use an organized, uniform computing architecture that easily scales up to hundreds of thousands of servers without much additional overhead [13].

Already in 2012, many hyperscale data center operators such as Facebook, Google and Apple have committed to using 100% renewable energy and many companies have followed (even though the Chinese giants such as Baidu, Alibaba and Tencent have made no such commitments and account for roughly 20% of the DC market). Google itself is the largest corporate buyer of renewable energy on the planet [13]. Yet, it remains to be seen if this 'hyperscale shift' will continue to offset the power consumption associated with the expected increase in data traffic. As the number of less efficient traditional DCs has already dropped substantially, it seems more likely that the global DC power consumption will increase in the next 5-10 years.

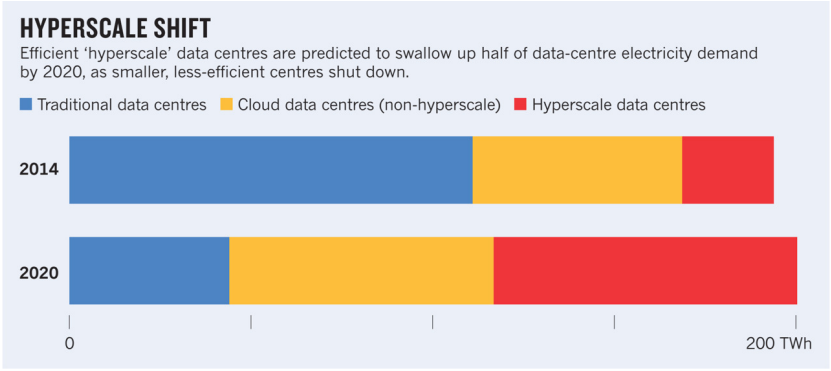

Figure 1.9: Power consumption distribution over types of data center: as much more efficient hyperscale DCs have replaced traditional DCs the total power consumption from 2014 to 2020 has remained relatively stable even though the DC workload has tripled (from [13]).

1.3.4 Towards Faster and More Efficient Interconnects

However, increasing the number of fiber links by building more and/or bigger data centers is not the only way in which the data processing capabilities can be increased. Another option is to expand the capacity per link by:

1. Increasing the speed by which a data packet is sent over a given link
2. Increasing the amount of data within a packet

Moreover, to sustain an estimated annual growth rate of 25 to 30% of the global internet traffic [4] both options are actively explored by the industry. This will require the development of compact and power-efficient optical transceivers capable of operating at double or quadruple the rate of the generation of DCIs at the beginning of this research (i.e. from 10-25 Gb/s per lane to 100 Gb/s and more).

In order to contribute to this research and development, the objective of this work is to investigate and design compact, low-cost, and low-power optical transceivers for next-generation intra and inter data center interconnects.

The following chapter will introduce some of the key aspects of optical fiber communication with respect to datacom applications. We will look at the main requirements, difficulties, and tradeoffs (both physical and economical) of optical interconnects in general and optical transceivers in specific, in order to provide context for the design choices of the fabricated optical transceivers and the experiment setups in the following chapters. As this research centers around integrated optical transceivers, the drawbacks and advantages of the two most prevalent platforms for the realization of photonic integrated circuits (silicon and indium-phosphide) will be discussed. Finally, an overview of the different types of high-speed modulators most commonly found on these platforms will be presented, as well as how these modulators can be used to realize multileveled modulation beyond the conventional direct conversion of the electrical signal to the optical domain on a single modulator.

References

[1] CERN, *Brief history of the internet*, [Online; accessed Dec. 2019], 1997. [Online]. Available: https://www.internetsociety.org/resources/doc/2017/brief-history-internet/.

[2] T. W. B. DataBank, *Individuals using the internet*, [Online; accessed Dec. 2019], 2018. [Online]. Available: https://data.worldbank.org/indicator/IT.NET.USER.ZS?end=2018&start=1960&type=shaded&view=chart.

[3] The World Bank: DataBank, *Worldwideweb: Proposal for a hypertext project*, [Online; accessed Dec. 2019]. [Online]. Available: https://www.w3.org/Proposal.html.

[4] Cisco, *Visual networking index: Forecast and trends 2017-2022*, [Online; accessed Sept. 2019].

[5] Sandvine, *2019 global internet phenomena report*, [Online; accessed April 2020], 2019. [Online]. Available: https://www.sandvine.com/press-releases/sandvine-releases-2019-global-internet-phenomena-report.

[6] H. R. Max Roser and E. Ortiz-Ospina, "Internet", *Our World in Data*, 2019, https://ourworldindata.org/internet.

[7] P. J. Winzer and D. T. Neilson, "From scaling disparities to integrated parallelism: A decathlon for a decade", *Journal of Lightwave Technology*, vol. 35, no. 5, pp. 1099–1115, Mar. 2017.

[8] Hugh McIntyre, *The Global Music Industry Hit $19 Billion In Sales In 2018, Rising By Almost 10%*, [Online; accessed April 2020], 2019. [Online]. Available: %7Bhttps://www.forbes.com/sites/hughmcintyre/2019/04/02/the-global-music-industry-hits-19-billion-in-sales-in-2018-jumping-by-almost-10/#2fc6df0418a9%7D.

[9] Global Market Insights, *Hyperscale data center market*, [Online; accessed April 2020], 2019. [Online]. Available: https://www.gminsights.com/industry-analysis/hyperscale-data-center-market.

[10] G. Zervas, H. Yuan, A. Saljoghei, Q. Chen, and V. Mishra, "Optically disaggregated data centers with minimal remote memory latency: Technologies, architectures, and resource allocation [invited]", *IEEE/OSA Journal of Optical Communications and Networking*, vol. 10, no. 2, A270–A285, Feb. 2018.

[11] Z. Zhang, Y. Deng, G. Min, J. Xie, L. T. Yang, and Y. Zhou, "Hsdc: A highly scalable data center network architecture for greater incremental scalability", *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 5, pp. 1105–1119, May 2019.

[12] G. M. Insights, *Google investeert 600 miljoen euro voor vierde datacenter in saint-ghislain*, [Online; accessed April 2020], 2019. [Online]. Available: %7Bhttps://www.datacenterknowledge.com/cloud/analysts-there-are-now-more-500-hyperscale-data-centers-world%7D.

[13] N. Jones, "The information factories", *Nature*, vol. 561, no. 7722, pp. 163–166, 2018.

[14] Y. Sverdlik, *There are now more than 500 hyperscale data centers in the world*, [Online; accessed April 2020], Oct. 2019. [Online]. Available: https://www.datacenterknowledge.com/cloud/analysts-there-are-now-more-500-hyperscale-data-centers-world.

[15] "Iso/iec/ieee international standard - information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements - part 3: Standard for ethernet amendment 10: Media access control parameters, physical layers, and management parameters for 200 gb/s and 400 gb/s operation", *ISO/IEC/IEEE 8802-3:2017/Amd.10:2019(E)*, pp. 1–376, Feb. 2019.

[16] Niall Robinson (ADVA Optical Networking), *Coherent vs. direct detection in metro data center interconnectivity*, [Online; accessed April 2020, 2017. [Online]. Available: https://www.datacenterknowledge.com/industry-perspectives/coherent-vs-direct-detection-metro-data-center-interconnectivity.

[17] John Houghton, *Understanding data center interconnect: Inter-data center*, [Online; accessed April 2020], 2018. [Online]. Available: https://www.neophotonics.com/what-is-inter-data-center-interconnect-dci/.

[18] C. Xie, "Optical interconnect technologies for datacenter networks", in *Datacenter connectivity technologies: principles and practice*, F. Chang, Ed., Stylus Publishing, LLC, 2019, ch. 1, pp. 1–32.

[19] Google, *Positive energy: Belgian site becomes first google data center to add on-site solar*, [Online; accessed April 2020]. [Online]. Available: https://sustainability.google/projects/belgium-solar/.

[20] International Energy Agency (IEA), *Tracking buildings*, [Online; accessed April 2020], 2019. [Online]. Available: https://www.iea.org/reports/tracking-buildings.

2

Optical Interconnects

2.1 Fiber-Optic Communication

2.1.1 From Copper to Fiber

Since the invention of the telegraph in the 1830s and 1840s, the amount of data and the distance over which it can be sent have increased enormously (Fig. 2.1). In the following century the growth of wired communication was mainly driven by transmission of electrical signals (voltages and currents) over copper wires. The early electrical communications systems used pairs of wires (e.g. the familiar twisted pair telephone wires). The first coax-cable (a metal core within a cylindrical metal cladding providing the return path for the current) was deployed during the Second World War and had a bandwidth of 3 MHz. However, at higher speeds the frequency-dependent cable losses rose rapidly, limiting the maximal bandwidth-length (BL) product. Hence, it was impossible to transmit over a large distance *and* at high data rates. Multi-gigabit transmissions over more than a kilometer only came within reach by switching to another physical medium: light.

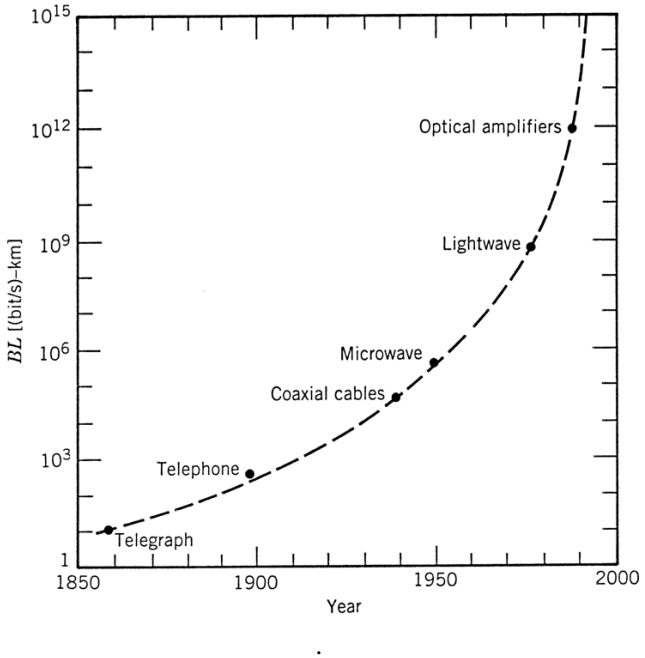

Figure 2.1: Rise in bitrate-distance product from 1850 to 2000 annotated with the emergence of key technological advancements [1]

2.1.2 Light as communication medium

As light is an electromagnetic wave, five physical dimensions form the basis for communication: polarization, frequency, quadrature (or phase), time and space. Fig. 2.2 shows some concrete examples of how each of these dimensions can be employed in the optical domain.

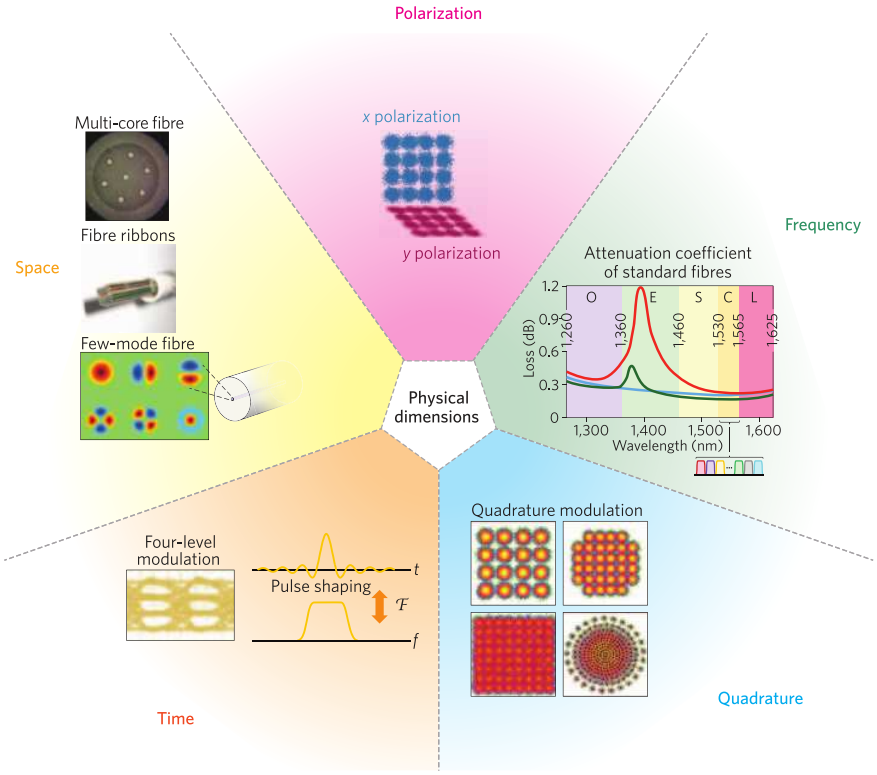

.

Figure 2.2: Examples of how the five physical dimensions of electromagnetic communication can be applied in optical communications [2]

Mathematically, we can represent light in an optical waveguide (e.g. a fiber) as:

$$\begin{pmatrix} E_x \\ E_y \end{pmatrix} = \begin{pmatrix} E_x(t)e^{j\omega t+\phi_x(t)} \\ E_y(t)e^{j\omega t+\phi_y(t)} \end{pmatrix} = \begin{pmatrix} (I_x(t)+jQ_x(t))e^{j\omega t} \\ (I_y(t)+jQ_y(t))e^{j\omega t} \end{pmatrix} \tag{2.1}$$

In most of the deployed optical communication systems, amplitude ($E_i(t)$, or as intensity: $|E_i(t)|^2$), phase ($\phi_i(t)$) and/or frequencies (ω) are used

to represent the information state of a transmitted symbol. The two orthogonal polarization states (x and y) and wavelength of the light wave are often exploited to increase the number of channels to carry theses symbols. These schemes are called polarization division multiplexing (PDM) and wavelength division multiplexing (WDM), respectively.

Virtually all intra data center optical interconnects currently rely on intensity modulation and direct detection (IMDD). Here, the amplitude (or rather the intensity $|E(t)|^2$) of a laser source is modulated by an electrical driver. This can be realized either by varying the bias current that feeds the laser, i.e. direct modulation, or by driving a modulator after the light source, i.e. external modulation. Direct modulation has the advantage that it requires fewer optical components and is therefore a more low-cost solution. These devices are, however, typically limited in bandwidth, as well as extinction ratio as there is a speed penalty in switching the laser off beyond its threshold current. Furthermore, due to the different states when transmitting high output powers and low output powers, the effective round-trip time in the cavity of the laser is also modulated. This introduces a frequency shift between the generated zeros (low power state) and the ones (high power state), i.e. chirp. In the presence of chromatic dispersion this chirp will introduce an additional distortion penalty on top of the chromatic distortion, reducing the link span even more.

The most common examples are vertical-cavity surface-emitting lasers (VCSELs) and direct-modulated lasers (DMLs). Semantically, VCSELs are of course also a type of DML, but a DML is typically understood to be a distributed-feedback (DFB) laser (e.g. as in [3, 4]). The main difference between both types of semiconductor lasers is the direction of light generation: DFB lasers are in-plane light emitters (Fig. 2.3), while VCSELs -as the name suggests- emit out-of-plane (Fig. 2.4). On the plus side, this means that a vendor can do quick wafer-level testing on VCSELs, but it makes them less practical light sources for optical ICs as their light directions mismatch.

Most VCSELs are multi-mode (MM) sources, i.e. there is more than one distinct peak in their output spectrum. However, single-mode (SM) VCSELs also exist. Usually, they are engineered towards the conventional single-mode telecom and datacom wavelengths of 1310 nm and 1550 nm, as low-cost alternatives for in-plane DMLs. These SM devices are not hindered by inter-modal dispersion but generally exhibit a lower bandwidth than their 850 nm MM counterparts. Nonetheless, in combination with the right electronics, state-of-the-art SM-VCSELs have

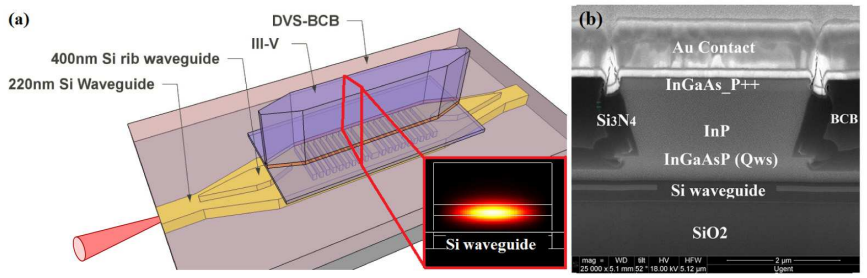

Figure 2.3: (a) Schematic of a IIIV-on-Si DFB laser capable of 28 Gb/s with the lasing mode predominantly confined in the IIIV-waveguide; (b) annotated cross section of the DFB laser [4].

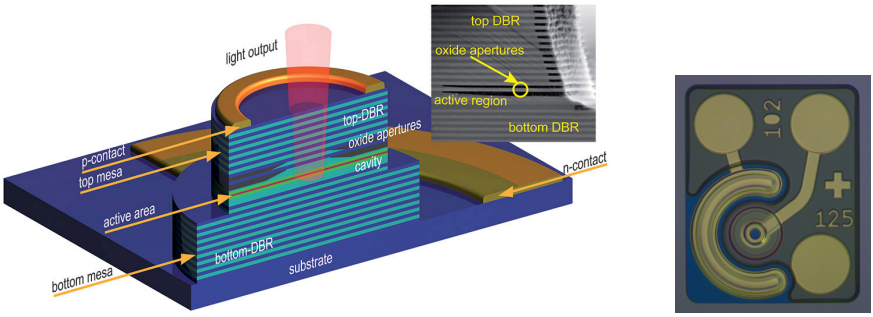

(a) **(b)**

Figure 2.4: (a) Cross section of a GaAs VCSEL [5]; (b) Top view of a 56 Gb/s PAM-4 datacom VCSEL [6]

demonstrated capabilities beyond 50 Gb/s with multilevel modulation in C-band [7] or with on-off keying over more than 15 km in O-band [8].

External modulators on the other hand can provide higher extinction ratios and higher bandwidths but require a more complex and costly optical frontend. By separating the modulation from the light generation, a laser source meant for continuous-wave (CW) operation can be optimized towards higher output power and efficiency without worrying about speed requirements. While directly modulated lasers are almost exclusively used for intensity modulation, external modulators are also widely used to create other types of modulation such as amplitude and phase modulation, or a combination of both.

2.1.3 Optical Fiber

Optical fibers are thin wires that consist of two different materials: the core and the cladding (that surrounds the core). Due to the difference in refractive index n between both materials ($n_{cladding} < n_{core}$), practically all of the light that is sent into the core will remain confined in it. Characteristically, fiber possesses extremely low losses per km (with the current record standing at 0.14 dB per km [9]). The loss profile of the most prevalent type of fiber made from silica glass, called standard single mode fiber (SSMF), is shown in Fig. 2.7. There are two clear transmission windows corresponding to the local minima in the loss curve: one around 1300 nm (the O-band) with a loss of 0.4 dB/ km and one around 1550 nm (the C-band) with a loss of 0.2 dB/ km. Unsurprisingly, these are the classic telecom wavelengths for optical communication. Especially C-band and L-band (around 1600nm), where the attenuation is the lowest, are the wavelengths of choice for long-haul optical communication systems as they allow for the minimal amount of inline optical amplifiers or repeaters, i.e. erbium-doped fiber amplifiers (EDFA) [10]. Invented in 1987 [11], EDFAs meant an big leap forward for optical interconnects as they provided an extremely broad gain medium, covering the whole C-band and parts of the L-band, with a tolerable amount of added noise. Consequently, several wavelengths could be easily amplified simultaneously, allowing to boost the fiber capacity drastically through wavelength-division multiplexing (WDM). In O-band, optical amplification is typically realized by semiconductor optical amplifiers (SOAs), but their performance is not comparable to that of EDFAs. SOAs have higher noise, provide less gain over a smaller wavelength range, and behave non-linearly with fast transit times. Hence, almost all optical links that require a decent amount of optical amplification are operated in C- and L-band.

2.1.4 Multi-Mode versus Single-Mode Fiber

There is another window around 850 nm (2.5 dB/ km or more) that is used for short-reach data communication up to a couple of hundred meters. This window is mainly attractive for use in datacom applications due low-cost optical components sources that can be realized at these wavelengths. However, to reduce cost even further these types of links typically use multi-mode mode fiber (MMF) which has a much wider core (50 to 65 µm) versus almost 9 µm in SMF (Fig. 2.5). This means that the fiber alignment to the MMF-components is much more relaxed tolerances, resulting in a significantly cheaper packaging. MMF-based optical transceivers almost exclusively use VCSELs as optical transmitter [12]. Thanks to the larger

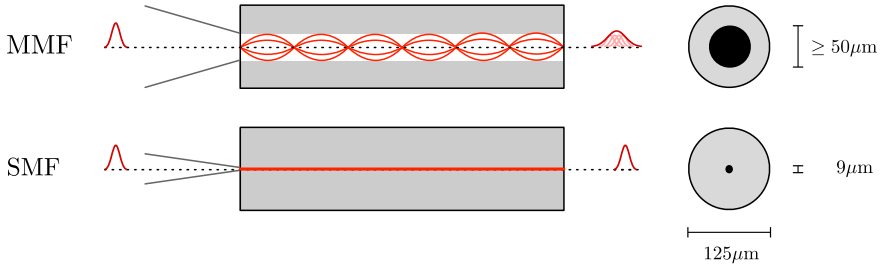

Figure 2.5: Representation of pulse propagation in multimode and single-mode fibers.

fiber core and thus more relaxed fiber alignment to the VCSEL, and the large number of short optical links in data centers, these low-cost transceivers have traditionally dominated the 100 m links in many data centers with unrivaled cost.

As the name suggest, MMF allows different modes of the light to be excited simultaneously due to the larger MMF waveguide core. However, because each mode travels at a different speed and thus a slightly different propagation delay, the modulated pulses will spread out in time and distort neighboring pulses, resulting in inter-symbol interference (ISI). In MMF the main source of dispersion arises from this *inter-modal dispersion*, limiting the useful distances to a couple of 100m depending on the MMF category. Currently, the fastest categories (OM4 and OM5) allow for a bandwidth-distance product of around 4.7 GHz · km (at 850 nm)[13].

Due to the more compact geometrical dimensions of the core, only one optical mode can be excited in SMF on the other hand. This eliminates the possibility of inter-modal dispersion and allows for much longer fiber spans. As the speed and the number of plus-100 m interconnects have increased significantly with the emergence of hyper scale data centers, SMF-based optical transceivers have seen their share grow significantly inside the data center as they offer a much higher bandwidth-distance product with respect to competing technologies such as MMF and the copper wireline connections. And this trend will only become more pronounced with increasing link rates (Fig.2.6).

However, even within a single mode there is a different propagation speed for each frequency component. This type of intra-modal dispersion is called *chromatic dispersion* (CD) or group-velocity dispersion. The amount of CD at a certain wavelength is given by the dispersion parameter D (expressed in ps/nm km) :

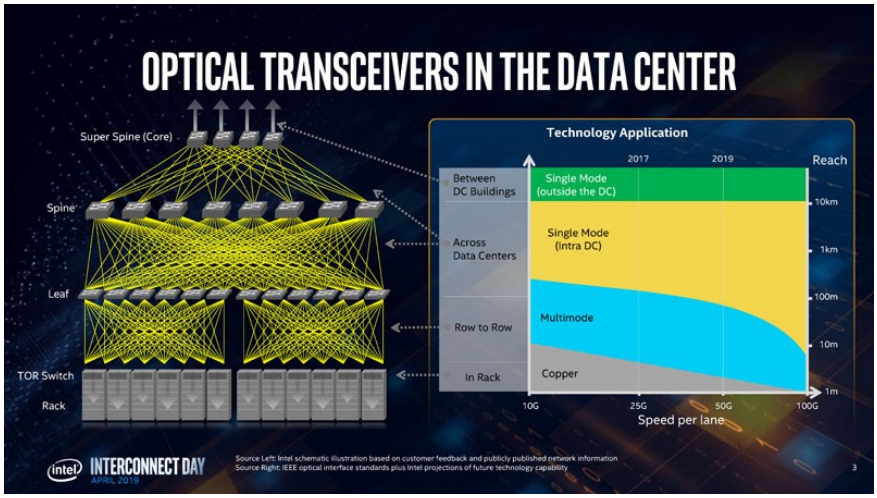

Figure 2.6: Evolution of different interconnect technologies and their deployment inside and between data centers according to Intel [14].

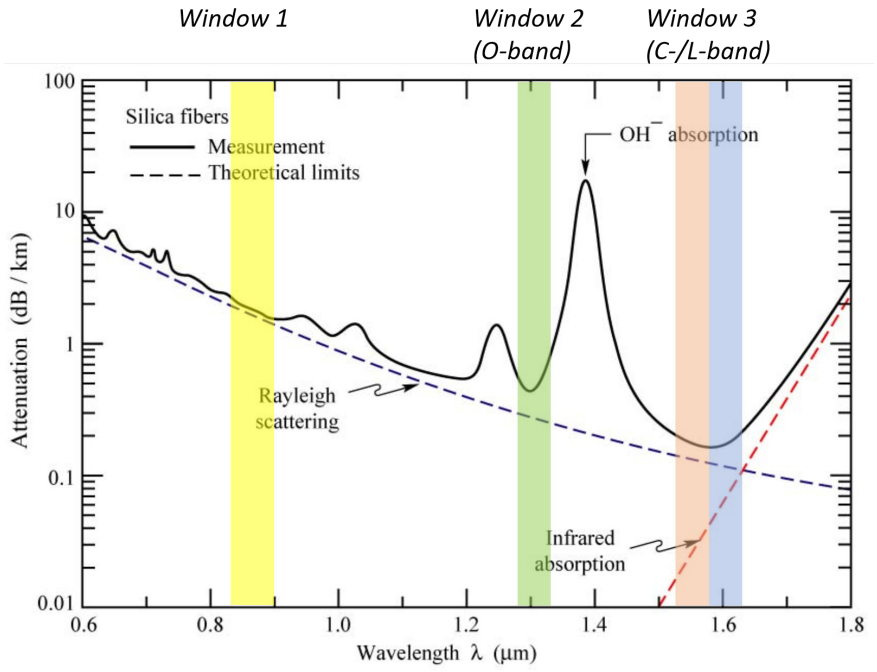

Figure 2.7: Loss profile of a standard SMF [15].

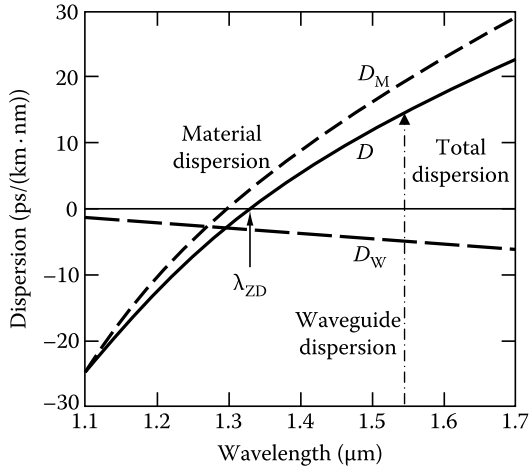

Figure 2.8: Chromatic fiber dispersion (D) consisting of material (D_M) and waveguide (D_W) dispersion of a standard SMF [16]

$$D = \frac{-2\pi c}{\lambda^2}\beta_2 \tag{2.2}$$

where β_2 is the group velocity parameter and λ the carrier wavelength in vacuum [17]. The evolution of D over wavelengths is shown in Fig. 2.8. It is worth to have a closer look into the effect of CD on the fiber transmission. In a small-signal regime the transfer function of a SSMF of length L can be approximated by:

$$H_{SMF}(\lambda_0, f) = e^{-\alpha(\lambda_0)L} \cdot e^{-j\pi D(\lambda_0)\lambda_0^2 L f^2} \tag{2.3}$$

with α the fiber attenuation, D the dispersion parameter and L the length of the fiber [1]. Ignoring fiber losses and assuming a chirpless source, we can describe the power transfer function of the fiber channel (e.g. after reception with a photodiode) as a filter with intensity variations related to the CD ([17]):

$$|H_{SMF}(\lambda_0, f)|^2 = \left|\cos \frac{LD\lambda_0^2\omega^2}{4\pi c}\right|^2 \tag{2.4}$$

with $\omega = 2\pi f$. This equation characterizes a frequency-selective channel with transmission nulls located at:

$$f_{notch} = \frac{2\pi}{\lambda_0}\sqrt{\frac{c(k+1/2)}{LD}} \tag{2.5}$$

with $k = 0, 1, 2, \ldots$. In Fig.2.9 the power transfer function $|H_{SMF}|^2$ is plotted for a transmission at 1550 nm (i.e. the center of C-band where the dispersion parameter D is approximately 17 ps/nm km) for the most common intra-DCI fiber lengths. We can see that for longer fiber spans (1) the number of transmission nulls increases in a given frequency window and (2) that the first notch appears at lower and lower frequencies. For a IMDD link the chromatic distortion penalty clearly limits the bitrate-distance product.

With the newest generation of optical interconnects in the data center operating at 53.125 Gbaud (e.g. 400GBASE-DR4 [18]), there is already a noticeable penalty after 1 km and transmission over more than 2 km will be challenging without any CD compensation for wavelengths in the C-band.

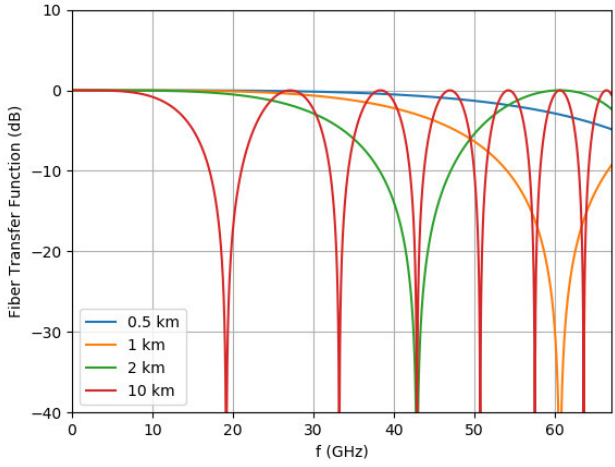

Figure 2.9: Transfer function of standard SMF at 1550 nm with $D = 17$ ps/nm km.

Interestingly, the dispersion curve crosses zero around 1310 nm (Fig. 2.8), resulting in significantly reduced chromatic distortion in O-band with respect to C-band. This also explains why O-band is currently the preferred wavelength window for SMF transceivers in data centers; the

links are short enough to avoid the use of optical amplifiers but they are becoming fast enough to feel the chromatic distortion in the C-band. As good rule-of-thumb, a 3- dB bandwidth of at least $2/3$ of the baudrate is required to ensure transmission with minimal impairments [19]. Looking at Fig.2.9, it is clear that from 40-50 Gbaud on it will be challenging to sustain a decent link performance up to 2 km of standard SMF for C-band transceivers without additional compensation techniques.

Nevertheless, in chapters 3 and 4 we will discuss two possible solutions where with the right combination of photonics and electronics transmission over more than 2 km was achieved in the C-band for rates from 50 to 100 Gbaud without any digital signal processing.

2.2 Photonic Integrated Circuits

As photonic integrated circuits (PICs) are growing more and more into a mature design domain, attention has slowly but surely started to shift from the design and optimization of pure building blocks (phase shifters, photodiodes, splitters and combiners,...) to the circuit level design. The designer needs to know the properties of the components he or she is using, but no longer necessarily has to know all the details of how a particular device is made in the fab. This evolution is very similar to the one electronics underwent in the last 30-40 years, where the characteristics of the device are boiled down to a handful of design parameters that can quickly capture the main performance of a device and which are fine-tuned and/or optimized with the help of EDA software tools.

Of course, this is only possible with a well characterized Process Design Kit or PDK: a library of the available building blocks in the given technology, where all components are listed with a -preferably small- performance variation interval to their nominal behavior, as well as a combination of layout and design rules to ensure this behavior over the complete wafer. Meaningful simulations on (very) large scale integrated (VLSI) circuits are only worthwhile when this kind of information is provided by the fab so that IC designers can verify if a circuit will achieve the desired performance with a sufficiently large confidence interval. Currently, integrated photonics is not at this point yet and most of the PICs will (have to) limit themselves to a relatively modest amount of building blocks and functionalities, where the designs still rely heavily on the experience of the photonic design engineer to anticipate these non-idealities and make the design robust to as many of them as possible.

Photonic Platform Choice: to Si or not to Si

Several photonic integration platforms exist, each with their own advantages and disadvantages. An important figure of merit for any integration platform is the feature size of building blocks. For PICs, the size of passive optical building blocks such as waveguides, splitters, filters, multiplexers, and demultiplexers are strongly correlated with the index contrast: the higher the difference in refractive index between the waveguide core and its cladding, the more confined the optical mode can be and the smaller the feature size necessary to perform a certain optical functionality. In Fig.2.10, the size of an arrayed waveguide grating (AWG), an device that can be used as a wavelength (de)multiplexer, is given for three popular integration platforms: Silica on Silicon, Indium Phosphide and Silicon on insulator (which is one of the most common forms of silicon photonics). The dependence of the feature size of the AWG on the index contrast is shown clearly in Fig. 2.10.

(a) Silica on silicon **(b)** Indium Phosphide **(c)** Silicon on insulator

Figure 2.10: Size comparison of an arrayed waveguide grating realized on (a) Silica, (b) InP, and (c) SiP. Thanks to the high index contrast SiP can realize much tighter bends (from [20]).

	Silica on silicon	Indium Phosphide	Silicon on insulator
Contrast	0.01-0.1	0.02-0.5	1.0-2.5
Mode diameter	8 µm	2 µm	0.4 µm
Bend radius	5 mm	0.5 mm	5 µm
Size	10 cm^2	10 mm^2	0.1 mm^2

Table 2.1: Size comparison of an arrayed waveguide grating.

However, as we will see, for high speed applications the active optical capabilities are equally important, and the choice of platform becomes dependent on many more parameters. For PIC designs geared towards

datacom applications, Indium Phosphide (InP) and Silicon photonics (SiPh) are the most prevalent choices as integration platform, and therefore deserve a more detailed discussion of their specific pros and cons.

Indium-Phosphide (InP)

Indium-Phosphide (InP) is a group III-V semiconductor material that is characterized by a direct bandgap and a very high electron velocity, making it a popular choice for ultra high-speed and high-frequency photonic and electronic integrated circuits (PICs and EICs). The direct bandgap of the III-V materials allows to make efficient lasers and photodiodes. The primary advantage of III-V based platforms is that by tweaking the bandgap of the materials through different composition, one can make efficient amplifiers, modulators, and photodetectors in the same platform. This means that all the building blocks are present for a basic transceiver. Unsurprisingly, the first successful PIC, dating back to 1987, was an InP-based electro-absorption-modulated laser (EML) [21].

Unfortunately, because it is scarce (Indium is about as common as silver on earth), InP is also quite expensive. Furthermore, the material is very brittle, making it nearly impossible to grow wafers much larger a couple of inches in diameter. In practice, wafers of 2 to 4 inches are most common. This severely limits the possibility of high volume production, which otherwise could be used to drive down the die cost. Another downside is the lack of a good native oxide. This means that the index contrast, needed to make low-loss optical waveguides with small bend radii and compact structures, has to be obtained by adding a different material, typically GaAs, to realize InGaAs or InGaAsP waveguides with InP cladding (or air at the edges).

Finally, the active-passive interfaces can be problematic as both the actives and the passives are made by epitaxy, introducing many interfaces and consequently many chances for defects. Moreover, many active components (lasers, amplifiers, modulators, detectors) actually require tailored material compositions and doping conditions for optimal performance. This is why the design of InP-based building blocks is often referred to as *bandgap engineering*.

So, to combine many of different functionalities on a single InP-PIC there is often a trade-off made between (1) shared material/process conditions that allow to implement all the required optical functions and (2) several material regrowths to tailor the bandgap to the functionality of each optical building block. Taking into account that if such a shared condition can be found it will always be a comprise in terms of performance between

all subsystems, while on the other hand every additional regrowth will significantly affect the yield (as the number of interfaces and the chances of defects grows), and increase the cost and complexity of the fabrication process. If a performance compromise is not tolerable, the integration of the different functions can of course also be realized by combining two or more different dies together (e.g. by butt coupling) at the expense of even higher complexity in terms of electrical/optical routing and packaging as well as overall cost.

Silicon Photonics (SiP)

On silicon on the other hand, the raw performance of (some of) the active optical components might be lagging behind those of their InP counterparts (or not present in case of optical gain), but much more functionalities can be combined in a single die. For example, polarization rotators and combiners are not readily available on InP, requiring an additional external subsystem to perform this task (e.g. in Silica). Many passive building blocks (routing, splitters, ...) can be made on a much smaller footprint, saving valuable transceiver real-estate.

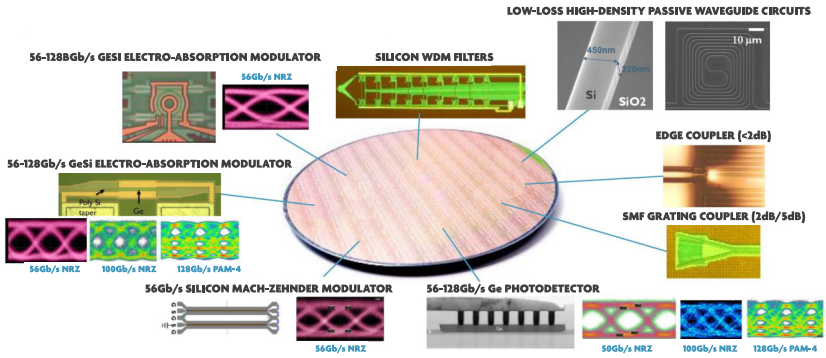

Figure 2.11: Main building blocks available on imec's 300 mm Silicon Photonics platform (iSiPP50G) [22]

Let us, for example, consider one of the most complex PICs: a long haul fiber transceiver. A silicon photonic version consists of only one PIC that realizes all the needed optical functionalities (modulation, detection, splitting/combining, routing, and polarization handling) apart from light generation which is accomplished by a co-packaged laser (e.g. [23]). Remarkably, the implementation with a external high-power, low-linewidth laser source might actually be more power efficient as only

the laser needs to be cooled, resulting in significant power savings on the thermal electrical controller. As the SiP PICs are less temperature sensitive than the InP PICs, close or even fully monolithic integration with electronics (driver, TIA, DSP) is much less of an issue as it is for InP [24].

Unfortunately, Silicon is a group IV semiconductor with an indirect bandgap. This immediately pinpoints to the main weakness of the silicon-on-insulator (SOI) platform: it lacks a native means of optical amplification, and as such, it cannot produce a light source. However, several solutions exist to alleviate this problem [25]: lasers can be grown epitaxially on the SOI (heterogeneous integration) or butt-coupled to the PIC during assembly without having to interfere with the SOI. Recently, transfer printing has gained quite some attention as a promising and -most importantly- cost-effective way to selectively transfer pieces of III-V material (e.g. to be used as gain medium for a laser) or even complete III-V devices to a SOI wafer or PIC [26].

Whereas indium is a very rare material, Si on the other hand is the most common material on earth (roughly 27% of the mass of the earth's crust is silicon). It is mechanically very strong and exhibits an excellent native oxide: SiO_2. As an oxidized Si wafer can be molecularly bonded to another Si oxidized wafer, nearly all waveguides in SiP can be made without any epitaxy, allowing for high PIC yields. Due to its mechanical strength, wafers up to 300 mm (12 inch) are commonly available for SiP PICs. Furthermore, these wafers can be fabricated reusing standard CMOS fabrication tools and processes, and benefit from 40 years of experience and investments in electronic IC fabrication. For a Si waveguide embedded in SiO_2, there is a high refractive index contrast between Si ($n \approx 3.5$) and SiO_2 ($n \approx 1.45$) both vertically and horizontally, while in InP high index contrast is limited to one dimension when contrasted to air. This allows ultra compact designs and wavelength scale integrated optical functions that are typically not readily available in InP such as polarization rotators and combiners, vertical fiber-to-chip grating couplers, and high-Q micro ring resonators (MRM) with bend radii 5 µm or less [27, 28].

2.3 Integrated Optical Modulators

As a large part of this work is focused on the development of new and advanced modulators, it would be instructive to briefly review the most common optical modulators and to understand their specific strong points and weaknesses. Furthermore, the interferometric operation of the Mach-Zehnder modulator is interesting to investigate in more detail, as this

principle will also used in the developed modulators. Direct modulation of a light source as is done with VCSELs or DFB lasers, where varying the current through the laser diode changes the optical output power, will not be discussed, as these devices are generally significantly slower than external modulators, produce frequency chirp, and are limited to power modulation schemes. Nevertheless, interesting hybrid approaches mixing direct and external modulation can be envisioned. For example in [29], an optical digital-to-analog converter (DAC) was generated by varying the current to a DFB laser as well as the taper from the laser to the waveguide (acting as an EAM) with two independent pseudo-random bit streams (PRBS) to create PAM-4.

Here, we give an short overview of the most common modulators in integrated photonic platforms: the Mach-Zehnder modulator (MZM), the ring modulator (RM), and the electro-absorption modulator (EAM).

2.3.1 Mach-Zehnder Modulator (MZM)

In a Mach-Zehnder interferometer (MZI), light is split in two paths where each path induces a different phase and then recombined again, resulting in either destructive or constructive interference depending on the accumulated relative phase difference between both arms. The induced phase difference can be fixed (e.g. using different optical path lengths) or it can be induced externally by applying an electrical signal to a phase modulator located in one or both arms. This structure is called a Mach-Zehnder Modulator (MZM).

Although only one phase shifter is required for amplitude modulation, MZMs are almost always made with a phase modulator in both arms as this provides several advantages. Fig. 2.12 shows the phasor diagram for an MZM with only one phase modulator which introduces chirp and a MZM with two phase modulators driven in push-pull. When driving both phase modulators differentially as illustrated in Fig. 2.12b, only half the voltage swing is needed per arm to induce the same level of modulation. This optical push-pull operation is often preferred as it makes the modulator inherently chirpless. This means that pure amplitude modulation is obtained without any undesired residual phase modulation. It also demonstrates that twice the phase shift (and thus twice the voltage swing) is required to produce the same modulation without the push-pull operation (Fig. 2.12a). A single MZM-arm needs to induce a total phase shift of π, while both phasors -when driven differentially- need only $\pi/2$ each: $+\pi/2$ on one arm and $-\pi/2$ on the other.

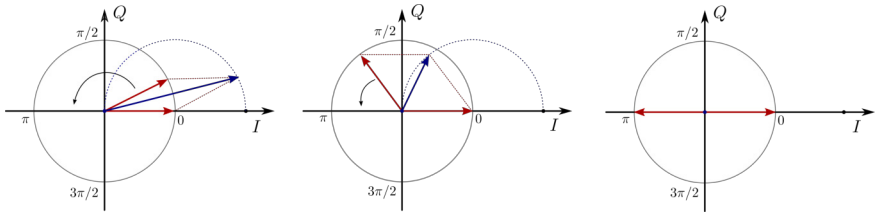

(a) MZM with single phase modulator. The phasor of the output field follows a half circle arc in the first quadrant of the complex plane. The resulting modulation has a varying phase component, i.e. the modulator chirps.

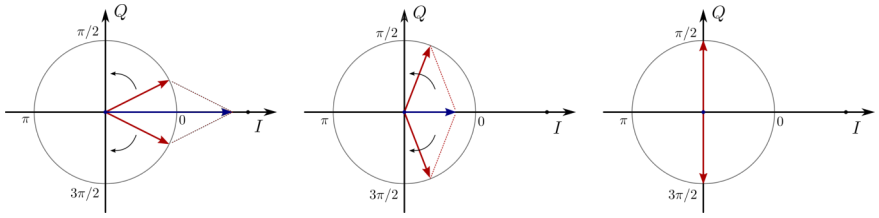

(b) Push-pull driven MZM. The phasor of the output field follows the real axis of the complex plane. The resulting modulation is thus chirpless as the phase is constant.

Figure 2.12: MZM operation: single phase modulator versus dual phase modulator in push-pull

By design a dual-drive MZM can produce several modulation mechanisms: phase, amplitude and power. For power modulation, the MZM is biased at 90° (quadrature), where DC-balanced swings will see the most linear region for power modulations such as NRZ and PAM-4. Setting the bias to 180° (minimal transmission) will result in the most linear region for amplitude modulation as shown in Fig. 2.13b. If both arms are driven with the same signal (and the same sign), pure phase modulation is achieved.

In an ideal phase modulator, the effective refractive index (Δn_{eff}) of the waveguide varies linearly with the applied voltage, which leads to a linear relation between the applied voltage and the induced phase difference as given by

$$\Delta\phi(V) = \phi_2(V) - \phi_1(V) = \Delta n_{eff}(V)\frac{2\pi L}{\lambda} \tag{2.6}$$

This linear electro-optic effect, called Pockels effect, is what is used in classic (LiNbO3) modulators which are able to generate high quality

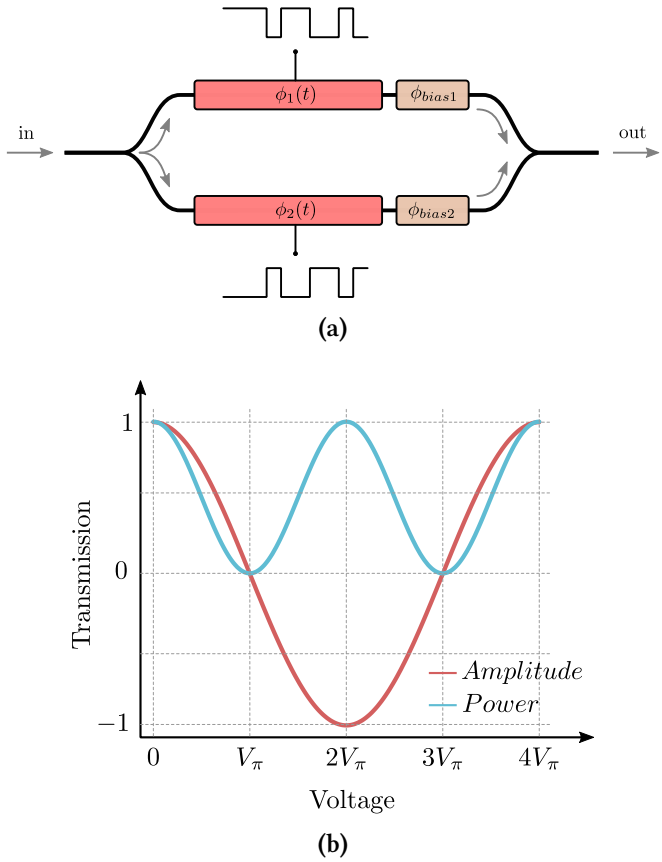

Figure 2.13: (a) schematic of a push-pull driven MZM; (b) transfer function of a MZM in function of the total phase difference $\Delta\phi$ between both arms.

pure amplitude modulation. Unfortunately, due to its centro-symmetric crystal structure, there is no Pockels effect present in silicon and thus the free-carrier plasma dispersion effect is exploited to create phase modulation [30]. Various modulation mechanisms using the plasma dispersion effect have been studied such as carrier depletion, carrier accumulation and carrier injection. All of them rely on essentially the same mechanism, i.e. by varying the distribution of free carriers around a waveguide, the optical mode feels a different effective refractive index (n_{eff}) which leads to a proportional phase change[31].

The rate at which n_{eff} changes with V is a material property, and is fully determined once an specific implementation and technology platform for

a phase shifter is chosen. This also leads to the main figure-of-merit (FoM) for MZMs: $V_\pi L$, typically expressed in $\mathrm{V}.mm$. It indicates how much voltage swing is needed to induce a 180° phase shift (V_π) for a given modulator length. It also acts as a useful design guide in determining the voltage swing for a given modulator length, vice versa, the minimal modulator length for a given swing to induce a certain amount of phase shift.

A popular implementation among plasma dispersion based phase modulators is the lateral PN-junction in depletion mode, as it is relatively easy to fabricate and tends to produce the fastest modulation. However, depletion introduces only a weak phase response leading to a high $V_\pi L$ (ca. 10 V mm to 20 V mm for state-of-the-art SiP MZMs). As a consequence, depletion type based phase shifters need to be very long (several millimeters) to accumulate sufficient modulation at feasible driver voltages. Hence, the electric contacts need to be driven as a transmission lines to account for the travelling wave effects, which in its turn requires the addition of a suitable termination (typically a resistor) to prevent reflections. Unfortunately, because the characteristic impedance of these electrodes as well as its terminating load is quite low ($\sim$30 Ω), the resulting current consumption in the driver IC is much higher when compared to compact modulators which can be driven as lumped components.

Another option would be to segment the electrode in sufficiently short pieces so that each segment can be driven as a lumped capacitive load by a distributed driver [32, 33]. Now, the resistive terminations can be avoided in favor of other low-power driver topologies (e.g. CMOS invertors). However, to keep a segment behaving as a lumped element with sufficiently low capacitance the maximal segment length is typically limited to 250 µm to 300 µm. Thus many segments and equally many drivers are needed for this distributed solution, possibly costing more power than one single transmission line driver and with a much more involved IO from the EIC to the PIC. The segmentation does offer other advantages as it allows the segments to be driven independently, which can be used to make multilevel optical modulation using only binary drivers (i.e. an optical DAC, more on this in Section 2.6) or to shape the frequency response of the modulator by applying delayed versions of the electrical signal to some segments and creating an optical feed-forward equalizer.

Either way, MZMs typically require significantly more power to be driven than small capacitive optical modulators such as RMs and EAMs.

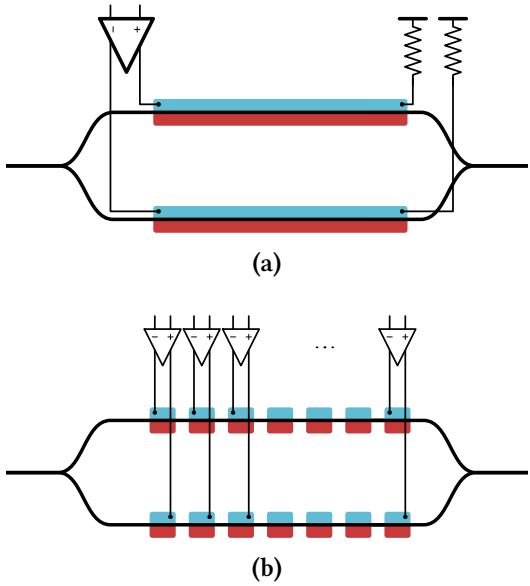

Figure 2.14: MZM topologies: (a) travelling wave electrodes of length L with a resistive termination driven as a transmission line; (b) N short, segmented electrodes (i.e. without noticeable traveling wave effects) of length L/N driven as lumped capacitive load by N drivers.

Velocity mismatch

Even without accounting for the more accurate non-linear dependence of Δn_{eff} on the applied voltage V in many integrated MZMs, Eq. 2.6 which led to design parameter and FoM of a constant $V_\pi L$, is only true for relatively "short" MZMs, as it assumes that the electrical and the optical wave travel at equal speed. This is however not the case. Somewhat counter-intuitively, on PICs the electrical waves typically travel faster than the optical waves through their respective medium, both in SiP and InP platforms. Hence after a certain propagation distance, there will be a noticeable walk-off between the electrical and the optical wave. This means that one has to actively slow down the electrical wave (e.g. by meandering the transmission line along the pn-junction [34] or adding active delay elements in the segmented case [32]) to compensate for this velocity mismatch. If not, this mismatch will manifest itself by a loss in modulation efficiency and, even worse, as a source of inter-symbol interference, limiting the EO bandwidth of the device. Although the amount of velocity mismatch is fixed for a given travelling wave

modulator, its negative effects become more pronounced at high symbol rates. For example, a walk-off of 2 ps might be negligible at 10 Gbaud, but it might be a clear source of interference at 53 Gbaud where the symbol period is less than 20 ps. Or in other words, as the MZM length is often chosen in terms of the available voltage swing, with rising baud rates velocity mismatch will have to be dealt with at shorter and shorter MZMs lengths.

2.3.2 Ring Modulator (RM)

Ring resonators have attracted a lot of research attention both as filters, sensors and modulators. They consist of a bus waveguide and a loop waveguide as shown in Fig. 2.15a. A fraction of the light is coupled from the bus to the loop, where it travels many times around the loop and the same fraction of the loop light is coupled back into the bus waveguide. For certain frequencies light coupling out will exhibit a π phase shift with respect to light in the bus waveguide and destructive interference occurs. The ring resonator is characterized by a typical periodically notched transmission spectrum where the free spectral range (FSR) between two neighboring notches depends on the circumference of the ring (or, equivalently, the round-trip time). Thus, RMs are a type of interferometer just as MZMs, where the the interference happens through feedback while in a MZM the interference takes place due to a recombination of the light in feedforward.

A ring resonator is transformed into a ring modulator by placing a phase shifter in the loop waveguide which can induce changes in the effective cavity length. This change shifts the resonance towards or away from the operational wavelength resulting in less (or more) light out the output of the bus waveguide, i.e. on-off keying modulation as shown in Fig. 2.15b.

Thanks to its intrinsically high index contrast, Silicon photonics is ideally suited to realize extremely compact ring resonators with high Q-factors and relatively high bandwidths, as these require low-loss waveguides with very small bend radii. Therefore, MRMs can be extremely compact modulators on Si-based platforms requiring only very short PN-junctions. Consequently, the associated junction capacitance is also very low, which allows for fast, low-power modulation, making the MRM well-suited for low-power capacitive-load drivers (such as CMOS invertors).

Recently, 128 Gb/s PAM-4 transmission was shown by Intel [35] using segmented MRMs with a bandwidth of 42 GHz. In [36], we demonstrated a 100 Gb/s PAM-4 single-segment MRM with more than 40 GHz bandwidth

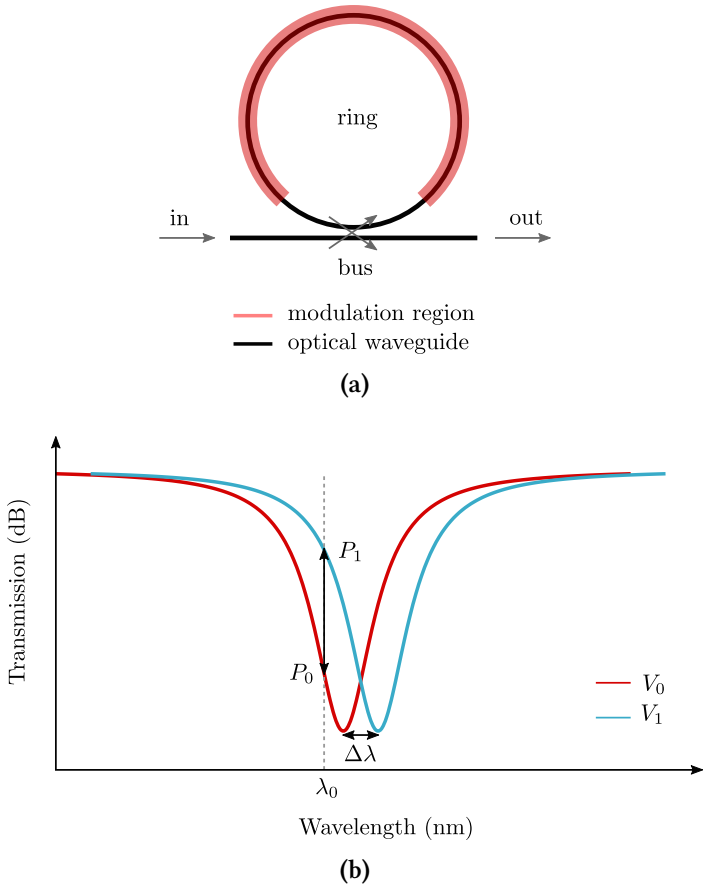

Figure 2.15: (a) Schematic and (b) transfer function of a ring modulator

in a joint experiment with imec, where the modulator was designed and fabricated.

Although the modulator itself does not consume much energy during operation, the peripheral circuitry to ensure proper operation might. Because RMs are very narrow band devices, their operational wavelength must be meticulously tuned to the laser wavelength and kept there. By design, they also suffer from a very high sensitivity to any type of disturbance such as small thermal or power fluctuations, changing the optimal wavelength of operation. This means that careful and dedicated monitoring and control circuitry must be provided in any real-life application, adding complexity, area and power consumption. These disadvantages only become more pronounced when using multiple MRMs that must be locked to the same wavelength or to a comb of wavelengths.

2.3.3 Electro-Absorption Modulator (EAM)

Both the MZM and the RM act as electro-refraction modulators, where a change in Δn_{eff} leads to a phase difference, which is subsequently converted into an amplitude modulation through interferometric recombination, either in feedback (RM) or in feedforward (MZM). Another way to generate amplitude modulation is by means of electro-absorption. Electro-absorption modulators (EAMs) exploit the fact that the bandgap of a semiconductor material will shift when a electrical field is present, changing the absorption spectra of the material. In bulk semiconductor materials this effect is called the Franz-Keldysh effect (FKE). A similar mechanism, called the quantum confined Stark effect (QCSE), exists in multiple quantum well (MQW) material stacks and is exploited extensively in IIIV-based EAMs.

Photons with sufficient energy will be absorbed and trigger the creation of an electro-hole pair, resulting in a photocurrent. For a given wavelength in the operational range of the modulator, an EAM can be thought of as a photodiode where the reverse bias regulates the fraction of photons that are being absorbed or, equivalently, the amount of the photocurrent that is being generated.

As the FKE is intrinsically a sub-picosecond process, EAMs in bulk materials such as Ge or GeSi as most SiP-based EAMs are typically RC-limited. Compared to conventional phase modulators, electro-absorption is a much stronger effect for a given interaction length with the optical waveguide. The change in absorption can be very significant at relatively low voltages, which means that EAMs can be made very short (typically

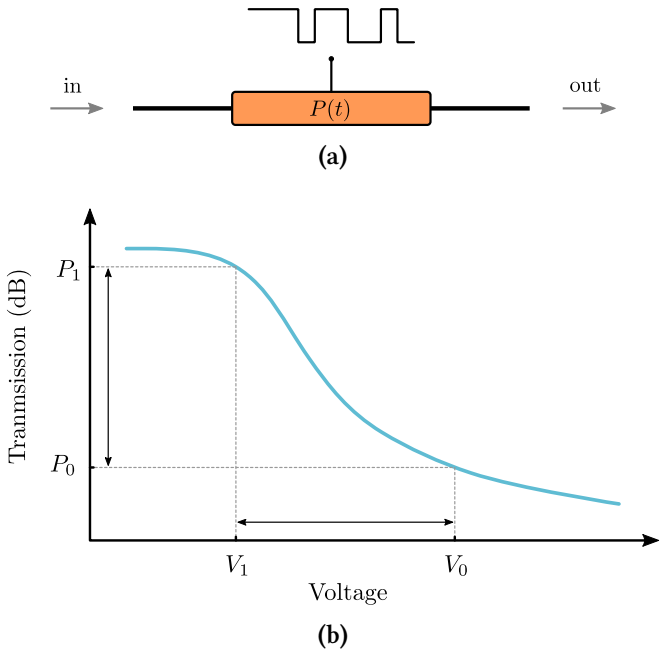

Figure 2.16: (a) Schematic and (a) transfer function of an electro-absorption modulator.

100 µm or less). Combining the sub-picosecond FKE and the low electrical parasitics associated with a very compact form factor, these devices can realize very high EO-bandwidths. For example, the standard GeSi-based FKE EAMs from imec are only 40 µm long and are capable of producing approximately 4 dB extinction at 2 V_{pp} with a EO-bandwidth well beyond 50 GHz. A higher extinction ratio can be obtained by increasing the voltage swing or by increasing the device length, at the cost of an increased optical insertion loss.

GeSi FKE EAM

In most of the transmitter designs in this work, longer versions of the GeSi EAMs developed by imec [37, 38] are used as basic building block for more complex modulators.

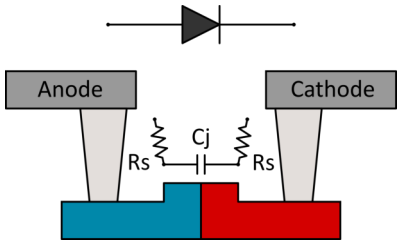

Figure 2.17: Cross-section of pn-junction (e.g. GeSi EAM), with a simplified equivalent circuit based on the series contact resistance (R_s) and the junction capacitance of the diode (C_j).

For realistic interaction lengths of this device, where we can still neglect travelling wave effects (<300 µm), the RC product is dominated by the junction capacitance and the contact resistance of the EAM [37]. Although these EAMs should be RC-limited, the measured bandwidth is barely affected by doubling the length from 40 µm to 80 µm as is illustrated in Fig. 2.18: doubling the length of the modulator doubles the capacitance (C_j), but halves the contact resistance (R_s). Using this RC equivalent circuit as a very rudimental model of the EAM, we can see that increasing the length should not drastically affect the bandwidth of the device. The bandwidth of the 80 µm long GeSi EAM was measured to be at least 67 GHz (Fig. 2.19). These observations correspond very well with the measurements done by imec on the 40 µm version available in their SiP-PDK which had only a drop of 1.5 dB at 50 GHz, hinting at a similar 3 dB-bandwidth as measured for the longer version.

A possible drawback of these EAMs is that their operational wavelength

range shifts to higher wavelengths when they heat up. However, this temperature induced drift can mitigated or kept within a certain range by preactively heating the EAM with resitive heaters. Based on the read-out of a nearby on-chip temperature sensor, the current through these heaters can be increased or decreased accordingly to ensure stable operation.

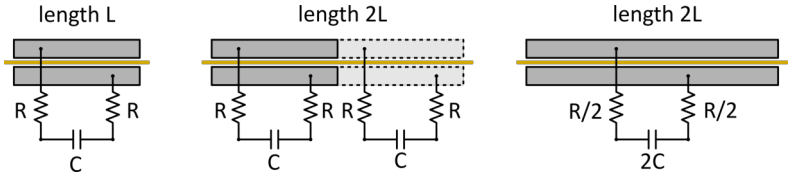

Figure 2.18: Based on the simplified RC-model from Fig. 2.17, the RC bandwidth of the modulator is not expected to change significantly if its interaction length is doubled.

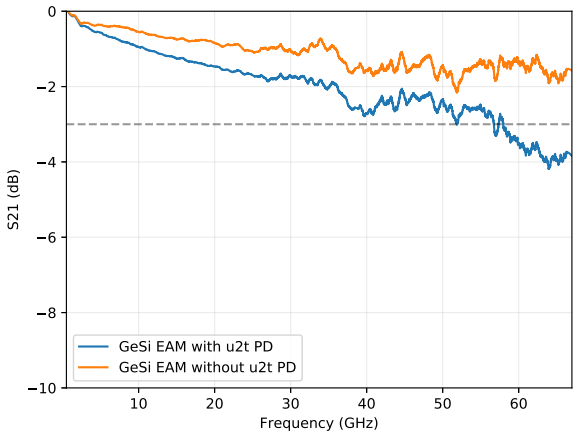

Figure 2.19: Measured EO-bandwidth of the 80 µm GeSi EAM (with and without de-embedding the frequency response of the 70 GHz photodiode).

But perhaps the main drawback of these Ge/Si Franz-Keldysh based devices is that it inherently lacks the capabilities to operate outside the C- and L-band. These are import bands for telecom but most datacom standards require operation in O-band where the chromatic dispersion is close to zero. Even so, we will show in this research that even in C-band 100 Gb/s/λ transmissions up to 2 km can be achieved at sufficiently low bit error rates with the right combination of electronics and optics, which covers all possible intra-data center links even in hyper scale data centers.

Nevertheless, much research has been devoted in the last years to realize multiple quantum-well based Ge/Si EAMs that are capable of O-band modulation as this would possible be a game-changer for silicon-based transceivers[39][1]. And although a silicon option might be preferred in terms of cost, QCSE EAMs made in IIIV-based materials such as InP are more than capable of efficient O-band modulation. For example, the EMLs used for 100 Gb/s/λ are de-facto DFB lasers with this type of EAM closely integrated next to the laser [41, 42]. Many of the topologies proposed in this work are therefore directly applicable to IIIV-based platforms or heterogeneous platforms where the III-V modulator is integrated on silicon photonics.

2.4 IMDD Optical Transceivers for Data Centers

Fig. 2.20 shows a typical gearboxed IMDD-based optical transceiver used inside data centers today. In this example, two streams of NRZ data with half symbol rate are multiplexed to one fullrate optical signal running as NRZ (at twice the input baudrate), or as PAM-4 (at the input baudrate but with four signal levels). Because the bitrate of the electrical channel from and to the optical module mismatches with the optical bitrate on the fiber, a so-called 'gearbox' is required to merge the data from two electrical lanes into a single optical lane. Either by multiplexing in the time domain (2×50 Gb/s NRZ to 100 Gb/s NRZ), in the amplitude domain (2×50 Gb/s NRZ to 50PAM-4), or possibly both (2×25 Gbaud PAM-4 to 50 Gbaud PAM-4). 400GBASE-DR4 is a prime example of a gearboxed link as it takes in 2×26.5625 Gbaud PAM-4 lanes to transmit one optical 53.125 Gbaud PAM-4 signal. Gearboxed solutions tend to add extra complexity and hence power consumption to the electrical design of the module as the both inputs need to be retimed and synchronized to each other before they can be joined. We will discuss the components in the gearbox version in more detail in the following section, as the non-gearbox version can be easily understood as a simplified version of this optical transceiver type.

2.4.1 Transmitter Side

At the transmitter side, multiple electrical input lanes are equalized, retimed by a clock-and-data-recovery circuit (CDR), and -in general- combined by a serializer (or multiplexer) into in a single electrical lane at a higher speed, although many alternative schemes exist depending on

[1]While writing this, the first monolithically integrated O-band EAM on a silicon platform was announced by Imec on VLSI 2020[40]

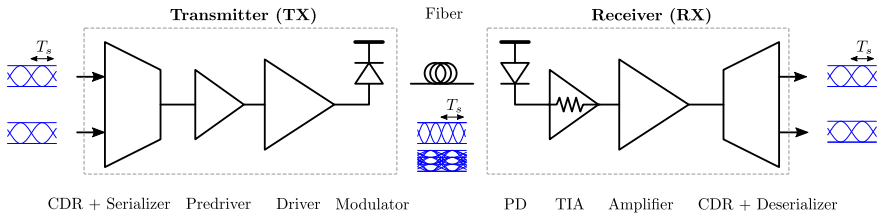

Figure 2.20: Simplified blockdiagram of a gearbox-based optical transceiver where two electrical NRZ lanes can be (de)serialized in time domain (NRZ over the fiber) or (un)stacked in the amplitude domain (PAM-4 over the fiber).

the specific configuration of the photonic components. A driver circuit then applies this signal with the appropriate voltage swing to one or more optical modulators.

With the transition from discrete optics to integrated photonics, many opportunities have risen to rethink the topology of drivers depending on the modulator's specific load impedance rather than a more generic 50 Ω interface (often accompanied by a transmission line if the modulator cannot be placed close enough to the electronic IC). Recent examples have demonstrated very high power efficiencies with less than 1 pJ/bit power consumptions for MRMs and EAMs [43–45]. Compact modulators such as MRMs and EAMs are inherently well-suited for these type of optimization as they behave as a small, lumped capacitive load, as explained in section 2.3. Segmenting the electrodes of a MZM allows similar driver topologies to be utilized on MZMs. But even though the per driver power consumption can be significantly lower, the cumulative power consumption of all segment drivers quickly surpasses that of more conventional resistive termination driver topologies.

A more interesting advantage of these segmented MZMs is the ability to generate multileveled optical signals based solely on binary electrical drivers. These types of optical DACs using multiple modulators either in series or in parallel have shown superior signal quality with respect to the stand-alone and multilevel-driven counterparts for MZMs [46], MRMs [47] and EAMs [48] as we will showcase in Chapter 4.

2.4.2 Receiver Side

At the receiver, the optical signal is converted to an electrical current by a photodiode (PD). A PD typically is a junction (often a p-i-n junction with

a small part of intrinsic semiconductor material between the p- and n-doped regions) where for a certain wavelength window (i.e. energy state of) the incident photons from an optical wave trigger the generation of an electron-hole pair, i.e. a current. The efficiency at which the conversion of optical power to current happens is called responsivity $\mathcal{R}$ (A/W):

$$\mathcal{R} = I_{ph}/P_{opt} \tag{2.7}$$

For optical powers below the saturation threshold of the PD, there is a linear relation between between P_{opt} and I_{ph}. This is the same mechanism that is exploited in EAMs, with the main difference that EAMs are operated at wavelengths close to their bandgap energy. This bandgap can be shifted by the presence of an electrical field (i.e. a voltage difference over the semiconductor junction) allowing for more or less photons to be absorbed. A photodiode, on the other hand, is designed to be operated further away from the material's bandgap to guarantee maximal absorption independently of the applied bias.

For high-speed communication, not only a high responsivity but also a very fast OE-conversion response time is desired. The response time of the detector is typically limited by (1) the transit-time i.e. the time it takes for the generated electron-hole pairs in the depletion region to reach the cathode and anode, and (2) the RC time constant of the internal equivalent electrical circuit of the diode. Hence, in high-speed applications PDs are reversed biased (typically between 1 V to 5 V) so that the depletion region can extend maximally, thereby reducing the junction capacitance and the resulting RC-product, and accelerating the electron-hole pair to the saturation velocity in the lattice. Contrary to EAMs this voltage does not affect the responsivity of the photodiode.

During the last five years integrated detectors have seen remarkable progress, especially in SiP, with state-of-the-art devices that can support up to 100 Gbaud operation (>67 GHz bandwidth) at high responsivities (>0.75 A/W), both in C-band [49, 50] and O-band [50].

However, another class of detectors exist where the electrical field over the junction is made so strong that an incident photon with the right frequency will introduce a chain reaction of electron-hole generations. The devices are called avalanche PDs or APDs and have a voltage dependent total responsivity due to the avalanche gain factor which multiplies the primary responsivity $\mathcal{R}$ of the PD with a factor $M(V)$:

$$I_{ph} = M(V)\mathcal{R}P_{opt} \tag{2.8}$$

In III-V based devices, very high bias voltages around 20 V to 25 V are not uncommon to ensure a strong enough electrical field to trigger the avalanche effect. In the past decade, many research has been dedicated towards Silicon-Germanium waveguide integrated APDs as they can provide much needed additional optical power budget for SiP links. Apart from enjoying the same advantages as SiP in terms of high-volume, low-cost and high-yield production, these SiGe devices can be operated at much more relaxed voltages. In [51], we demonstrated that a sub-5 V SiGe APD designed by imec [52] could support links up to 50 Gb/s PAM-4 with an average avalanche gain of 6 dB. Sub-5 V operation is an especially desirable target as it would allow the APD to be biased by the power supply of the transceiver module or even directly from the TIA without the need for dedicated off-chip high-voltage supply provided by a boost-converter.

Unsurprisingly, this improved responsivity comes at a price. As the avalanche process takes time to build up, APDs tend to be noticeably slower than their non-avalanche operated counterparts. Nevertheless, due to the inherent power budget advantages they could provide in many optical links, high-speed waveguide-integrated APDs are still a very active research domain [52, 53].

Leaving the optical domain, the photocurrent from the PD is converted to a voltage by a transimpedance amplifier (TIA). This signal is then equalized, retimed, and deserialized into multiple electrical lanes to be finally sent off-module.

From a system design point of view, the optical receiver leaves less room for different architectures or topologies: almost all IMDD transceivers use a single photodetector to perform the conversion to the electrical domain. As more and more PDs combine very good linearity, high bandwidth, and responsivities quite close to the quantum limit, it is difficult to improve upon a single PD implementation without quickly sacrificing one or more of these characteristics. The transmitter, on the other hand, offers several possibilities to differentiate and cater to the specific strengths of certain technologies, both optically and electrically: What kind of (external) modulator: ring, EAM, or MZM? Single modulator or multiple modulators? And if multiple, in series or in parallel? Driven binary or multileveled? Is the modulator resistively terminated, with or without transmission line, or can it be driven as a lumped capacitor? For this reason, most of the research activities in this dissertation focus on the transmitter in an attempt to demonstrate the benefits that might arise out of a joint electronics and photonics aware design.

2.4.3 Pluggable Optical Modules

Fig. 2.21 shows a some of the most common data center optical transceivers in pluggable modules. For 200G and 400G, the most popular industry standard form factors are the Quad Small Form Factor Pluggable Double Density (QSFP-DD) and the Octal Small Form Factor Pluggable (OSFP) module. The OSFP module is slightly larger than the QSFP-DD, making it less compact but this allows for a higher power consumption as it can more easily evacuate heat from the ICs. These form factors are separately defined by Multi-Source Agreement (MSA) bodies, specifying compact mechanical modules that are connectorized and pluggable into a compatible socket in a system platform (e.g. a switch). As different implementations are favored by different switch suppliers, pluggable module vendors typically cannot afford to limit their development to a single form factor and neglect a large part of the transceiver market.

	SR	DR	FR	LR	ER	ZR
Distances	100 m-300 m	500 m	2 km	10 km	40 km	80 km

Table 2.2: Inter and intra DCI fiber span nomenclature. SRx are MMF-based, all the other categories use SMF.

Depending on the requirements of the link, there is a wide variety of pluggable modules. Table 2.2 gives the different optical transceivers categories for ascending fiber reaches. Apart from the shortest distances (SRx), which is still the exclusive domain of cheaper MMF transceivers, all other categories employ SMF.

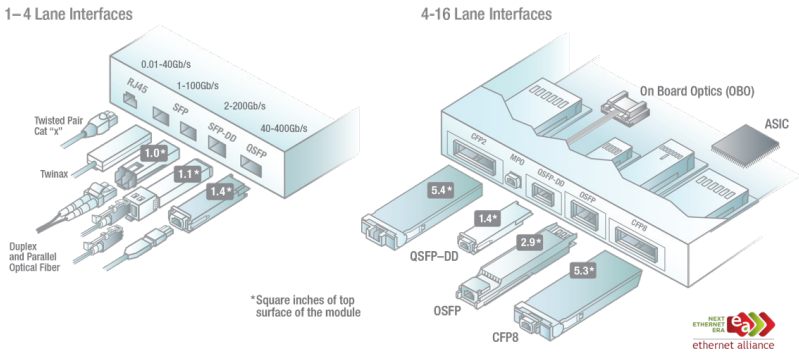

Figure 2.21: Common form factors for high-speed pluggable modules.

Another way to categorize the modules is based on their lane rate, both optically and electrically. Faster components can achieve higher baud rates, which reduces the number of optical lanes needed to reach a certain bitrate. Less optical lanes means less optical components and therefore a possibly cheaper solution. That is, if the electrical and optical components can support these higher baud rates without too much additional complexity and cost. Furthermore, a higher number of optical lanes does not necessarily imply a higher number of optical fibers. Due to the high carrier frequency associated with the traditional telecom and datacom wavelengths, the occupied bandwidth of even the fastest available optical transceivers is still three orders of magnitude smaller than the frequency of the laser source.

	400GBASE-SR16	**400GBASE-DR4**	**400GBASE-FR8**	**400GBASE-LR8**
Distance	100 m	500 m	2 km	10 km
Bitrate / lane	26.5625 Gb/s	106.25 Gb/s	53.125 Gb/s	53.125 Gb/s
Modulation	NRZ	PAM-4	PAM-4	PAM-4
Media	MMF	SMF	SMF	SMF
Number of lanes	16 lanes	4 lanes	8 lanes	8 lanes
Number of fibers	2×16	2×4	2×1 (WDM)	2×1 (WDM)

Table 2.3: Overview of the main properties of the 400G optical ethernet standards for DCIs.

For example, a laser at 1550 nm produces a light wave that oscillates at around 192 GHz, while a 50 Gbaud signal has less than 100 GHz of significantly contributing frequency components and can thus generously fit in a 100 GHz window (that is: 50 GHz left and 50 GHz right of the carrier). In C-band, 100 GHz constitutes to approximately 0.8 nm. Hence, C-band alone (1530-1565 nm) can carry up to 44 100 GHz wide channels. And a similar reasoning can be made for the other standard transmission windows. This mechanism of combining as many as possible optical lanes in a certain transmission band is called dense wavelength division multiplexing (DWDM). DWDM is used extensively in long-haul and submarine fiber links which are extremely expensive (hundreds of millions of dollars [54]) to deploy. Evidently, it is much more cost-efficient to squeeze every last *bit* of useful bandwidth out of a single point-to-point link than to add new links.

In data centers, however, the trade-off between multiplexing in space versus in wavelength is more balanced and both options have their place. For example, 100 Gb/s transceivers (carrying 4 optical lanes of 25 Gb/s) can

be found in two types:

- 100G-CWDM4: 1×2 fibers carrying each 4 wavelengths
- 100G-PSM4: 4×2 fibers carrying each 1 wavelength

CWDM4 is a coarse wavelength division multiplexing setup (CWDM) which combines four wavelengths in a single fiber, i.e. one fiber for the upstream and one fiber for the downstream link (Fig. 2.22b). PSM4, on the other hand, has four pairs of up and down stream fibers each carrying one wavelength (Fig. 2.22a). PSM4 requires more fiber alignments which is still a costly part of the packaging as the alignment tolerances are on the order of 1 µm. However, PSM-based modules can drop the optical multiplexer (mux) at the TX and demultiplexer (demux) at the RX that is needed in the CWDM setup. To relax the design constraints on the optical muxes and demuxes, the CWDMs have a much wider transmission band surrounded by even wider guard bands to ensures successful operation over large enough temperature range where the center wavelengths of both the muxes and the lasers can drift. In a PSM-based link the laser can drift much more, relaxing the constraints on the laser even further. As the optical lanes are physically separated on different fibers, a single laser is used as a shared light source for all optical channels. This makes the PSM ideal as a break-out module where a single PSM4 module can be connected to 4 other modules.

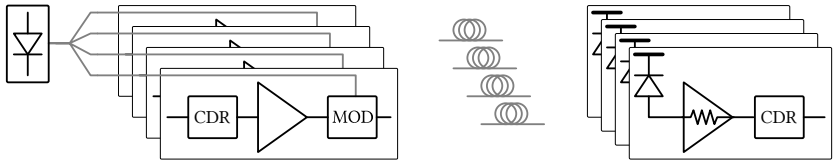

(a) PSM4 link

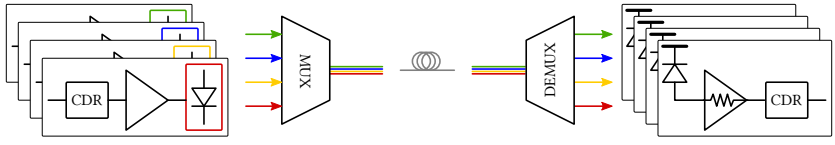

(b) CWDM4 link

Figure 2.22: Schematic of (a) a PSM4 and a (b) CWDM4 transceiver based interconnect.

As the optical lane rates increase, so do the electrical interconnects going from and to the optical modules. Running these differential traces over

several inches of PCB on the server is becoming an even bigger bottleneck than the optical IOs. As there is no room to deserialize the lanes further on the already crowded server PCB (with high-speed connector vendors now proposing some creative fly-over cable alternatives [55, 56]), the only practical solution is to increase the lane rates. This becomes problematic as these electrical channels endure significant frequency dependent losses. Therefore, the electrical IOs were among the first to consider different modulation formats than NRZ such as duobinary and PAM-4.

2.5 Intensity modulation formats: NRZ, EDB and PAM-4

Until recently almost all IMDD links inside the data center used simple but robust on-off keying (OOK) with only two possible symbols in the constellation: a low-power optical level representing a 'zero' and a high-power optical level representing a 'one', similar to bits in digital link. By far the most common implementation of this scheme is called, somewhat intuitively, Non-Return-to-Zero to indicate that the transition from one symbol to the next did not require drop in signal power contrary to the aptly named Return-to-Zero (RZ) where each symbol or unit interval (UI) is distinctively recognizable as it is starts and ends in the zero-state. Although RZ makes it intrinsically easier to perform a CDR operation due to the strong presence of the clock frequency in the signal's spectrum, the electrical and optical components have to be able to operate at twice the bandwidth of the equivalent NRZ signal. For RZ, the transmission of a '1' requires at least a rise time and fall time to fit in a single UI, while for NRZ only half a rise and fall time are needed.This means that for a given system bandwidth, NRZ can achieve twice the bitrate of RZ at the cost of a more complex clock-recovery.

For a long time NRZ has been the ideal fit for low-cost and low-power high speed optical interconnects, as it allows to combine efficient and low-complexity binary electronics.

However, the only way to increase the bitrate with these binary formats is to increase the speed at which the modulator is operated with the same amount. As a rule-of-thumb for an NRZ signal, a bandwidth of 60 to 70% of the desired bitrate is often deemed sufficient to carry a certain speed. For example, to successfully generate 50 Gb/s a modulator should have at least 30-35 GHz of bandwidth without needing equalization techniques. For a 100 Gb/s link this would already be 60-70 GHz, bandwidths that only

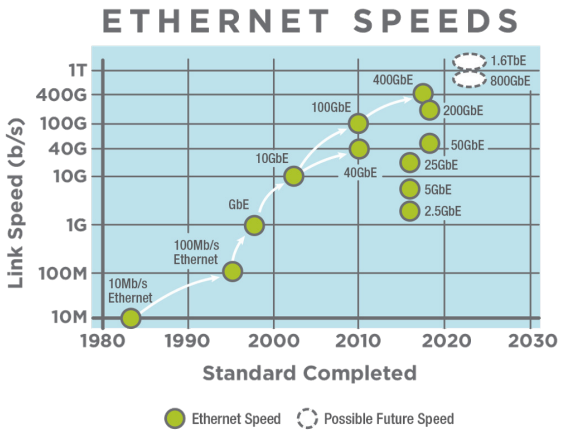

Figure 2.23: Evolution of the standardized gigabit ethernet speeds.

a few state-of-the-art modulators could support at the time and which are certainly not commonly available commercially.

Hence, to reach the 100 Gb/s target on a single optical lane set forth by the roadmaps for next-generation data center interconnects (Fig. 2.23), modulation formats that were more spectrally efficient than NRZ gained much attention. In particular, electrical duobinary (EDB) and 4-level pulse amplitude modulation (PAM-4) were investigated closely as they provided a decent trade-off between reduced bandwidth requirements and added complexity.

Although it took longer than expected, by the beginning of 2018 the newly adopted IEEE standard 802.3bs decided on PAM-4 as the modulation format of choice for the next-generation data center transceivers [57]. Likely, because this would require the least effort on the optical side to implement with the current generation of optics and because electrical I/Os to and from the pluggable modules had also recently adopted PAM-4 as the frequency dependent losses of the PCB traces became too difficult to overcome at 50 Gb/s NRZ and beyond.

Fig. 2.24 compares the effect of the each format (i.e. NRZ, EDB or PAM-4) on the time and frequency domain properties of the signal. It becomes clear that both formats trade in eye-height (i.e. sensitivity) in favor of more relaxed bandwidth demands.

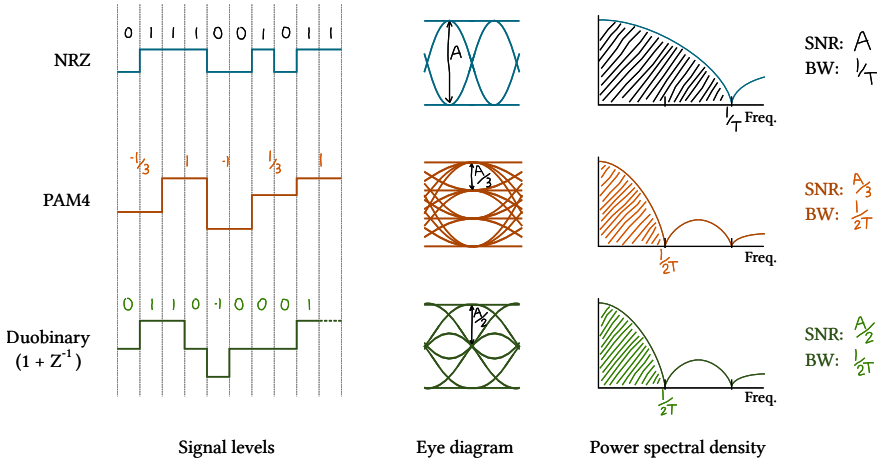

Figure 2.24: Comparison of time and frequency domain properties of NRZ, EDB, and PAM-4.

2.5.1 NRZ

NRZ is a very robust binary modulation format that can be transmitted and received with low-complexity and efficient non-linear electronics. It does, however, require the highest lane speeds for both the electronics and photonics. This means that it is also the least spectrally efficient format of these three, making it the most vulnerable for chromatic distortion.

2.5.2 EDB

EDB is closely related to NRZ. It requires only a relatively simple coder and decoder consisting of a XOR-port and a delay element and a conversion from a binary to a 3-levelled signal by a (1) a delay-and-add or (2) a low-pass filter resembling a 4th order bessel filter as illustrated in Fig. 2.25. In particular, the low-pass filtered version is very appealing as the limited frequency response of the TX, RX, or both can act as the EDB pulse shaping filter. Equalization at the TX or the RX can be used to further aid the pulse shaping, if the channel deviates too much from the ideal filter response.

Low-pass filtering leads to a 3-level eye diagram made up out of two characteristically pie slice shaped eyes, as depicted in Fig. 2.24. Hence, this 3-level signal needs to be sampled in two positions after reception, increasing the receiver's complexity. An example of such a receiver, consisting of two comparators and a XOR-port, is shown in Fig. 2.26. Fortunately, next to the relaxed bandwidth specifications, the input of

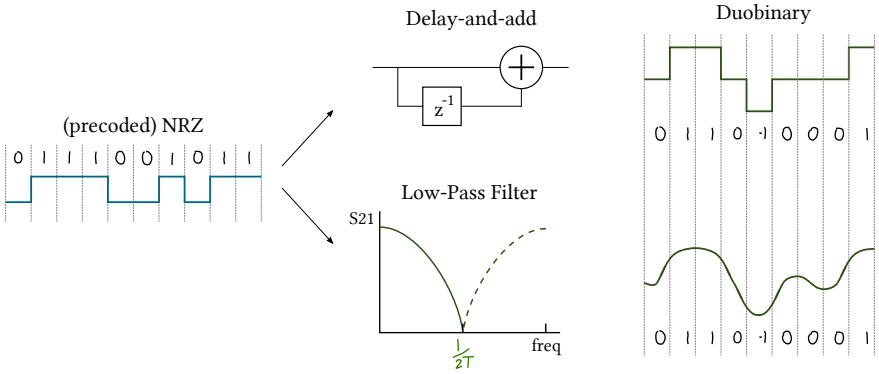

Figure 2.25: Electrical duobinary pulse shaping options: (1) delay-and-add or (2) low-pass filtering (approximating a 4th order Bessel).

the RX only needs modest linearity as the signal can tolerate quite some compression as long as it manifests itself symmetrically on the eye (Fig. 2.27). For many electronic ICs -and especially differential ones- this is indeed the case.

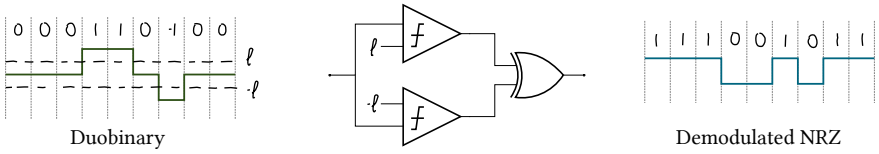

Figure 2.26: Example of an EDB receiver.

Additionally, the reduced bandwidth improves the link tolerance against chromatic distortion. Nevertheless, EDB still requires some of the most critical electronics (serializer, CDR, deserializer, ...) to run at full-rate, albeit with a lower bandwidth requirements. For the current Ethernet standards this would entail an operation at 106 Gb/s which might be too challenging for many of the current generation IC technologies, driving up the power consumption and/or cost of the IC.

2.5.3 PAM-4

PAM-4 has four levels and, hence, four symbols (0, 1, 2, 3) in its constellation diagram, each representing two bits (00, 01, 10, 11)[2].

[2]For non-gray coded data. For gray coded data the LSB_{gray} = XOR(MSB, LSB), MSB_{gray} = MSB resulting in (00, 01, 11, 10)

Therefor, this format can operate at half the symbol rate of binary formats and deliver the same data throughput. Furthermore, assuming the rule-of-thumb bandwidth reasonings of NRZ still holds, PAM-4 should also approximately halve the bandwidth requirements on the fastest components in the link.

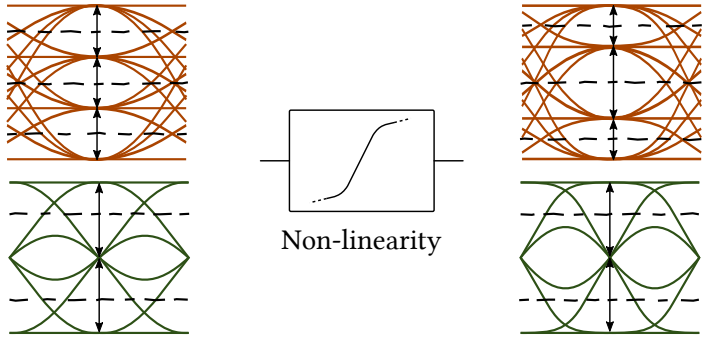

Figure 2.27: Comparison of the effect of (symmetrical) non-linearity on a EDB and PAM-4 eye.

However, there is no such thing as a free lunch. As PAM-4 scales in the amplitude or vertical domain (i.e. more levels) rather than the time or horizontal domain (i.e. more UIs per second), it places stronger and new demands on the transceiver circuitry in other areas, such as linearity and sensitivity. Linearity can be translated into the requirement that the four levels need to be uniformly distributed over the signal swing. If this is not the case, some of the eyes (typically, the outer ones) will be compressed with respect to the other(s) as illustrated in Fig. 2.27.

Even in ideal circumstances, PAM-4 introduces a minimal sensitivity penalty of 4.7 dB as each eye in the constellation is a factor 3 smaller with respect to a single NRZ eye at the same optical modulation amplitude (OMA, defined as $OMA = P_{high} - P_{low}$ with P_{high} and P_{low} the maximal and minimal eye levels in each constellation). In reality, this penalty can easily rise to 6 dB or more in the presence of inter-symbol interference (ISI) [58]. Non-idealities such as ISI are more pronounced on PAM-4 as there are twice as many possible transitions from any given symbol to the next. This also means that on average less crossings occur per eye, making the CDR operation more challenging. An implementation of a PAM-4 receiver is shown in Fig. 2.28, consisting of three comparators and two XOR-ports.

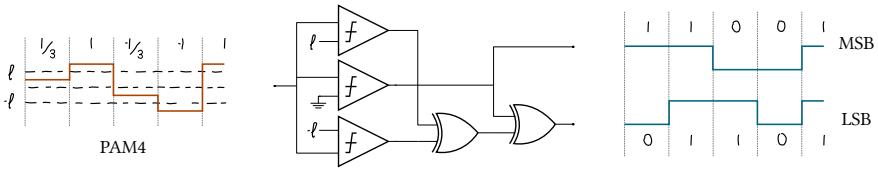

Figure 2.28: Example of an PAM-4 receiver and decoder which outputs the MSB and the LSB.

PAM-4 with non-linear modulators

Looking at the EO-characteristics of the modulators in Section 2.3, it is clear that all optical modulators offer a limited linearity at best. So even if the produced electronic symbols are evenly spaced PAM-4 symbols, the resulting optical symbols might be spaced unevenly. Fig.2.29 shows the effect on a non-linear electrical-to-optical conversion by the modulator and how they can be mitigated, at the cost of extra complexity or a reduced modulation depth.

A common compensation technique is to predistort the electrical levels in anticipation of non-linearity introduced by the modulator. Although this can be quite effective in guaranteeing equidistant signal levels, the non-linearity has not disappeared, it is merely masked by the compensation technique. E.g. linear equalizers, commonly used to extend the bandwidth, will be less effective than with more linear devices.

Furthermore, additional circuitry needs to be added to implement this compensation technique. If the compression by the modulator is symmetrical, as is the case for MZMs, the LSB could be simply given higher amplitude (or the MSB lower) before both are combined in an analog version of the driver. However if the compression is not symmetrical, as with MRMs or EAMs, the implementation becomes much less straightforward. Now the levels need to be set independently from each other. In a DSP-based implementation, the electrical DAC would require significantly higher number of effective bits (ENOB) to ensure sufficient resolution to place the electrical levels.

This is why another technique, which simply limits the electrical modulation to the most linear region of the modulator, is also used even though it typically leads to dramatically lower extinction ratios with respect to a binary driven device. Of course, this option assumes that there still is a sufficiently linear region with sufficient extinction ratio (at least 3.5 dB [57]). Many integrated modulators are intentionally operated at

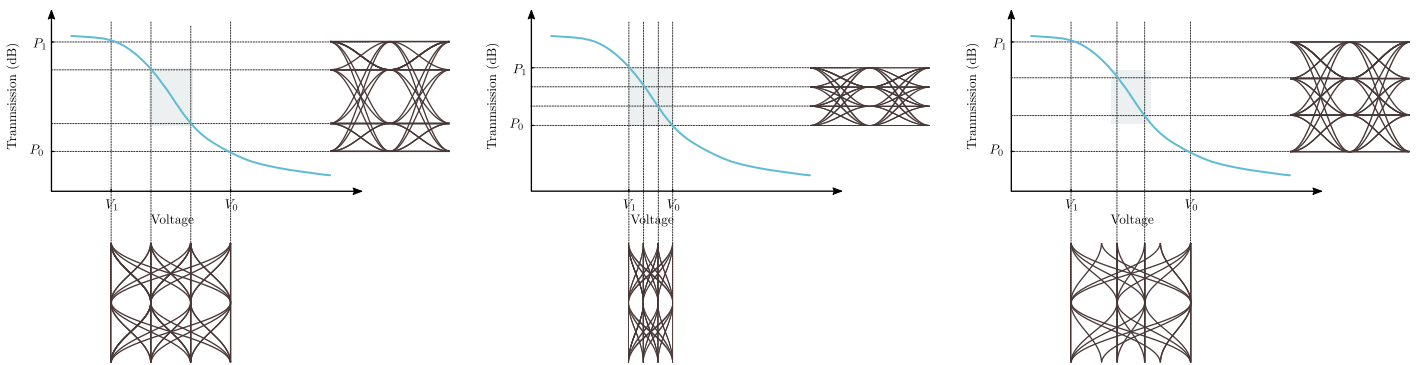

Figure 2.29: Effect of a non-linear EO-conversion on a PAM-4 signal: (a) without compensation, (b) after limiting the swing to the most linear region at the cost of extinction ratio, and (c) after predistorting the drive levels with the same amount as compression introduced by the optical modulator.

much lower extinction ratios than they can provide. This is because higher ERs either require stronger drivers with higher electrical swings and thus more power consumption, or larger modulators with longer interaction lengths, and thus lower bandwidths, higher insertion losses, and larger footprint.

Finally, no matter the specific electronics architecture, inherent to the non-linear nature of the problem, the ideal relative location of a level changes with the swing. This means that a 'set-and-forget' approach is often not desirable, or even feasible. Keeping the relatively high temperature ranges and the component variability in mind, some feedback on the real-time optical signal is almost always preferred.

2.6 Optical Digital-to-Analog Converters

Although multilevel modulation provides several benefits by increasing the spectral efficiency (i.e. the number of bits per second per nanometer optical bandwidth), it introduces new design challenges as well. Section 2.5 already illustrated some examples of how this would influence an electrical design, especially in terms of linearity. However, this is a problem that does not necessarily has to be solved in the electrical domain. Depending on the implementation, performing the DAC operation in the optical domain can have several advantages over an electrical DAC. Taking the non-linear or even non-symmetrical transfer function of the most common optical modulators (cfr. Section 2.3) into consideration, it makes sense to look for topologies where these modulators can remain binary driven and yet still be able to produce a multilevel optical signal.

The idea behind many of these optical DACs (ODACs) is to adapt the optical components towards electronics, rather than the other way around. By balancing the advantages and disadvantages of each domain more evenly, binary-driven multi-modulator ODACs can show significant performance improvements over regular electrical DACs driving a single modulator.

In the past years several types of optical DACs have been proposed, and they can generally be divided into two groups:

- the modulation builds up *sequentially*
- the modulation builds up *in parallel*

In sequential modulation (Fig. 2.30a), the modulation is divided over several sections of modulators each connected by a single waveguide and

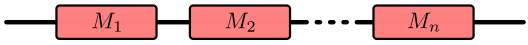

(a) Sequential modulation

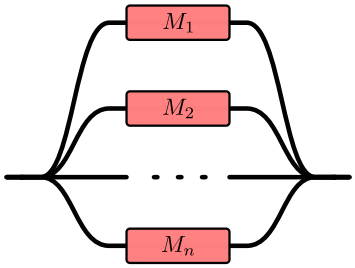

(b) Parallel modulation

Figure 2.30: Base topologies for optical DACs.

the modulation builds up as the light traverses each modulator. Examples are the segmented MZM and segmented MRM, where phase shifter is cut in several smaller pieces.

In parallel modulation (Fig. 2.30b), the light in each modulator has its own waveguide which modulates a fraction of the incident light and the multilevel modulation only arises after recombination as shown in Fig. 2.30b. The optical PAM-4 modulator using parallel EAMs developed during this research is a prime example of this method and will be discussed much more in depth in chapters 4 and 5.

In sequential modulators, the insertion loss of each individual modulator adds to the total insertion.

$$IL_{seq} = IL_{M1} + \ldots + IL_{Mn} \tag{2.9}$$

For a parallel modulator the overall insertion loss equals that of a single modulator.

$$IL_{para} = IL_{Mi} \quad \text{for} \quad i = 1, \ldots, n \tag{2.10}$$

That is, *if* all parallel branches are recombined in phase.

Another difference can be found in the time of modulation: in sequential modulation, light is modulated at different time stamps as it takes time for the optical wave to propagate through each modulator section. This implies that, if the total modulation length grows beyond a certain size, there will be a non-negligible time difference between the time when the

same portion of the light reached first and the last modulation section, i.e. velocity mismatch (Section 2.3.1). Hence, a compensation mechanism has to be implemented on the electrical side, e.g. by adding active or passive time delays, to realign the electrical wave fronts with the optical wave packet. This issue does not present itself in parallel modulators, as long as a layout can be found that ensures the equal electrical path length to each modulator.

Although all multi-modulator designs require a more involved contacting between drivers and modulators, the placement of the IOs of a serialized modulator can be still relative straightforward as the modulator can run parallel to the chip edge of the PIC and the EIC (for example, Fig.2.32). For an equivalent parallel implementation, a more careful layout of the optical modulators and the waveguides is typically required. Here, the advantage of the SiP platform comes into play, as it allows the PIC designer to route the waveguides in such a way that a similar contacting scheme to the chip edge can be achieved with minimal additional chip area or propagation losses in the waveguides. However, this also becomes cumbersome and complex once the number of modulators becomes too high.

In both case, however, the IO is limited by the available chip edge, which makes combining a multi-channel chip topology with a multi-driver modulator topology impractical for much more than two differential drivers per channel. For a 4-channel IC with two differential drivers per channel, this would already constitute at least 16 bondpads. If more modulators need to be driven, the multi-channel driver topology is typically abandoned in favor of multiple single-channel drivers (cfr. Fig. 2.32)

Of course, for fully monolithic platforms where photonics and electronics share the same chip real-estate, this might be less of an issue as the modulator IO does not need to adhere to a chip edge.

To provide a background for the pros and cons of the parallelized ODACs in the later chapters, the following two sections will elaborate on the two most common examples of serialized ODACs, i.e. the segmented MZM and the segmented MRM.

2.6.1 Segmented MZM

The most common serial optical DAC is probably the segmented MZM, sometimes also referred to as segmented electrode MZM (SE-MZM) or multi-electrode MZM. Section 2.3.1 explained how partitioning the long

phase shifter(s) can be utilized to prevent velocity mismatch by delaying the electrical signal to each section of the modulator to match the speed of the slower travelling optical wave. Instead of driving each segment with delayed copies of the same signal, independent drivers could be used which would produce an independent partial modulation with an amplitude proportional to the size of the segment with respect to the total modulator length. This way the characteristic cosine transfer function of a classic MZM could be sampled with a resolution depending on the number of segments, realizing 2^N levels with N the number of segments.

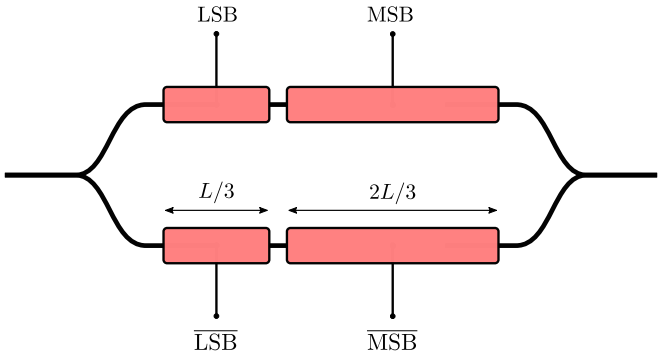

(a) Binary weighted

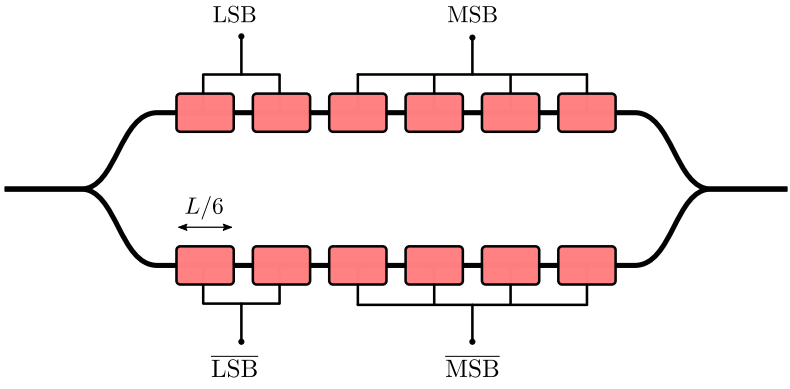

(b) Equally weighted

Figure 2.31: Example of two different implementations of a (push-pull driven) segmented MZM for PAM-4 modulation

As with any type of DAC, the contributions can be binary weighted, equally weighted, or a combination of both. An implementation example of both options for a PAM-4 modulator can be seen in Fig. 2.31.

In binary weighted SE-MZMs, each segment doubles in length which allows the bits to be mapped directly to a single specific segment. Conceptually, this is a very simple scheme that requires the least amount of segments -and therefor- the least amount of drivers. However, ideally the drivers should be independently optimized as they each steer a different (typically, capacitive) load, which increases complexity and design time. Furthermore, it might be difficult to ensure a sufficiently similiar performance for all the segments for DACs with a modest to high resolution (e.g. $N \geq 4$), because the difference between the smallest and largest load becomes too large to be fully compensated by the individual fine-tuning of the drivers.

In equally weighted SE-MZMs, each segment has the same length and the binary doubling is achieved by grouping several unit segments together. Typically, these topologies have much higher number of segments for a given modulator length. Once the unit segment becomes short enough to be treated as a lumped element, each segment can be driven as a small capacitor. This matches very well with advanced CMOS nodes where the analog performance suffers from low supplies (<1V), making it challenging to design high-swing resistive drivers, but where a relatively straightforward inverter based topology can be very effective to drive the MZM.

Although a single driver can be made quite power efficient as it does not need to be linear or drive a large load, the total number of drivers can be high and hence the cumulative power consumption can be too. Another disadvantage is the copious amount of I/Os between the electrical and the optical components associated with having many segments, which might pose a severe drawback due to the large number of bondpads.

This is one of the areas where a fully-monolithic technology combining (Bi)CMOS electronics and Silicon Photonics could potentially have a big edge over competing solutions, and hence has attracted quite some attention both from industry and academia [59–61]. Unfortunately, the number of foundries currently offering access to such a process is still limited [62, 63], but this could change in the near future as more and more of these fully monolithic electronic/photonic ICs (EPICs) find their way into products [61, 64].

Nevertheless, a binary-weighted implementation can be an especially attractive option for PAM-4 modulation as this would require only one cut to divide the phase shifter in two pieces of length $L/3$ and $2L/3$, corresponding to the LSB and MSB. Even though segments would still be too

large to be driven lumped -forcing a transmission line approach and hence additional power to drive the resistive termination- this implementation is often more power efficient and lower in complexity than the equally weighted implementation as the alternative would require driving shorter but also many more segments, likely swaying the power budget in the advantage of the binary weighted solution.

A drawback of both implementations is that they do not eliminate the non-linearity of the MZM, but rather offer new possibilities to compensate them. To compensate for the sinusoidal compression of the transfer function, one still needs to either predistort the effective modulation either by adjusting the physical length of the electrodes or by allowing several smaller segments to be regrouped in a different configuration. However, this requires additional complexity or fixes the DAC operation to a specific modulation format. As the amount of phase modulation per segment is proportional to the swing and the interaction length ($V_\pi L$ is a constant value), this compensation can also be realized in the electrical domain by providing each segment with a tunable electrical swing. This might be the easiest solution if the output stage of the driver is a differential pair, but it might not be trivial with an inverter-based design which typically offers a lower power consumption.

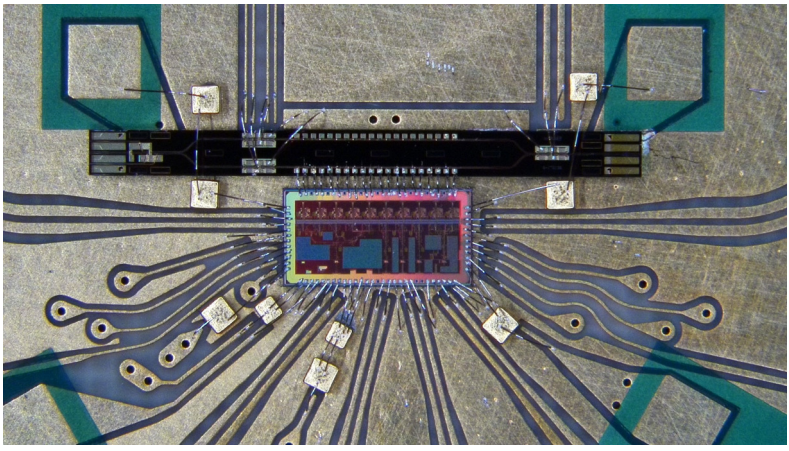

Figure 2.32: Assembly of a InP segmented MZM and a 10-channel CMOS driver capable or generating 25 Gbaud 8-ASK. For IQ modulation, another copy driver IC would be needed above the InP PIC.

Fig. 2.32 shows an example of an InP segmented MZM driven by a 40 nm CMOS IC made to generate an 8 level modulation (8-ASK) by using a

combination of both binary and equally weighted segments [32]. SE-MZMs with a high number of segments are mainly found in IQ-modulators for coherent optical communication as they are able to generate high order advanced modulation formats such as 32-QAM or 64-QAM with high quality[34]. It is clear that these multi-segmented ODACs require significantly more IOs between the PIC and EIC. This forces the driver IC to be parallel to the PIC and thereby prevents a single (2-channel) driver IC to drive both I and Q MZMs, which would be an option for an MZM-based IQ-modulator driven by a 2-channel electrical DAC. All of this adds additional size and complexity to the co-packaging of electronics and optics.

2.6.2 Segmented MRM

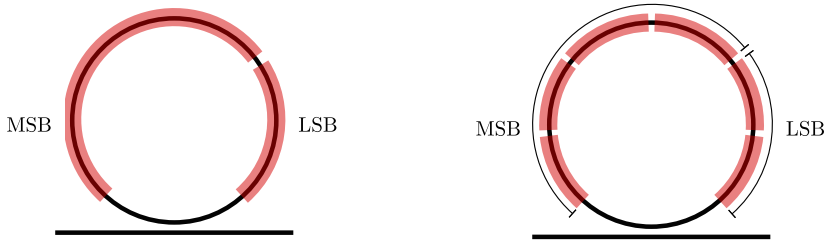

(a) Binary weighted **(b)** Equally weighted

Figure 2.33: Example of two different implementations of a segmented MRM for PAM-4 modulation

Another example of a serial ODAC can be found in the segmented microring modulator as shown in Fig. 2.33. As the MRM operation is also based on phase shifters, they share many common characteristics. However, as MRMs are extremely compact modulators, they face the opposite problem in terms of IO: because the bondpads are several magnitudes larger than the modulator itself, the very compact circular layout makes it difficult to provide a high number of I/Os associated with higher order ODACs. Therefore the MRM is probably best suited for 2-bit ODACs geared towards PAM-4 modulation. Again, a fully monolithic electro-optic technology could be advantageous as it would negate the need for bondpads between the modulators and the drivers all together (a route that is recently being explored by Ayar Labs [60, 61]).

These compact devices can be very fast and extremely efficient in terms of dynamic power consumption thanks to small modulation regions which can be driven as lumped capacitors. There is, however, also a significant static power consumption associated with the active monitoring

and control circuitry needed to track and stabilize the resonance over temperature and incident power as ring modulators are very susceptible to any kind of disturbance.

And just as with the segmented MZM and other serialized optical DACs, the non-linearity still needs to be alleviated by tweaking the effective interaction lengths (pre-silicon production) or the driving voltages per segment (post-silicon production).

2.7 Demarcation of the research domain

Based on the findings in this chapter, we can say that (1) external modulation is more interesting than direct modulation in order to obtain sufficiently high baudrates, (2) Silicon Photonics can offer low-cost optical transceivers in high volume and with high yield by using existing CMOS foundries, and (3) among integrated modulators EAMs have demonstrated an interesting set of properties: very high bandwidth, small footprint and low pJ/bit energy consumption at low drive voltages without the need for additional control circuitry to align operational wavelength to the laser. Thanks to their small footprint and low electrical parasitics, these modulators can be regarded as small, lumped capacitors which allows to use power efficient drivers, especially for binary operation. These properties make them attractive devices to be used in novel, power-efficient optical DAC topologies where, instead of a single multilevel-driven modulator, there are multiple binary-driven modulators.

For these reasons, this research focuses on the realization of compact and low-power Silicon transceivers based on EAMs for inter and intra data center connectivity. The following chapters will demonstrate several ways how these transceivers can be used to achieve 50 to 100 Gbaud operation per lambda.

2.8 Outline of the dissertation

This chapter provided a background for this work and introduced some of the key aspects of optical fiber communication with an emphasis on the ones that come into play for data center optical interconnects.

Chapter 3 demonstrates how 100 Gb/s NRZ modulation can be achieved with a single 80 µm long waveguide-integrated GeSi EAM in combination with a BiCMOS transmitter and receiver chipset. A follow-up experiment investigates how duobinary modulation could be used to improve the

resilience against chromatic distortion to extend the link reach beyond 2 km, covering all possible intra data center links. These were the first 100 Gbaud experiments with a silicon photonic modulator without requiring any DSP.

In chapter 4 we investigate how an intensity modulator with a non-linear and non-symmetrical transfer function such as a EAM can be used to generate PAM-4, the modulation format adopted by IEEE for the 200G and 400G Ethernet interconnects. This results in the development of a novel type of optical DAC where two binary driven intensity modulators are placed in an interferometer in quadrature (i.e. with a constant 90° phase difference).

In the following chapter, this new modulator concept is combined with the chipset from chapter 3 to demonstrate the first real-time silicon PAM-4 link above 50 Gbaud. We also elaborate on the advantages and disadvantages of inducing the power difference between the LSB and the MSB electrically or optically.

Next, we look beyond the (by now) current generation of optical transceivers running at 100 Gb/s per lane and try to tackle some of the obstacles for the next-generation transceivers which will run at double or quadruple the current lane rates.

Chapter 6 illustrates how an optical serializer can leverage current -or even previous- generation electrical and optical components running at 26 Gbaud to achieve 104 Gbaud lane rates with NRZ and PAM-4, resulting in the (then) highest IMDD bitrate on a silicon platform, without DSP.

Another path towards the future of optical transceivers is explored in Chapter 7, where we investigate how the design considerations set forth in the previous chapters can be used to realize compact and low-power coherent transmitters and receivers using EAMs.

Finally, in chapter 8, the main results of this work are briefly summarized together with some concluding remarks. An outlook towards the future implementations of data center transceivers wraps up the book.

References

[1] G. P. Agrawal, *Fiber-Optic Communication Systems*, 3rd ed. Hoboken, N.J.: Wiley, 2013, OCLC: 864892021. [Online]. Available: http://rbdigital.oneclickdigital.com (visited on 07/17/2018).

[2] P. J. Winzer, "Making Spatial Multiplexing a Reality", *Nature Photonics*, vol. 8, no. 5, pp. 345–348, May 2014.

[3] A. Abbasi, S. Keyvaninia, J. Verbist, X. Yin, J. Bauwelinck, F. Lelarge, *et al.*, "43 Gb/s NRZ-OOK Direct Modulation of a Heterogeneously Integrated InP/Si DFB Laser", *Journal of Lightwave Technology*, vol. 35, no. 6, pp. 1235–1240, Mar. 2017.

[4] A. Abbasi, J. Verbist, J. V. Kerrebrouck, F. Lelarge, G.-H. Duan, X. Yin, *et al.*, "28 Gb/s Direct Modulation Heterogeneously Integrated C-band InP/SOI DFB Laser", *Opt. Express*, vol. 23, no. 20, pp. 26 479–26 485, Oct. 2015. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-23-20-26479.

[5] H. Li, P. Wolf, P. Moser, G. Larisch, J. Lott, and D. Bimberg, "Vertical-cavity surface-emitting lasers for optical interconnects", *SPIE Newsroom*, Nov. 2014.

[6] II-VI press report, *II-VI Incorporated Introduces 56 Gb/s PAM4 VCSEL Arrays for 400 Gigabit Ethernet Transmission*, 2017 (accessed: 01.16.2020). [Online]. Available: https://www.ii-vi.com/news/ii-vi-incorporated-introduces-56-gb-s-pam4-vcsel-arrays-for-400-gigabit-ethernet-transmission/.

[7] W. Soenen, R. Vaernewyck, X. Yin, S. Spiga, M. Amann, G. Van Steenberge, *et al.*, "56 Gb/s PAM-4 Driver IC for Long-Wavelength VCSEL Transmitters", in *ECOC 2016; 42nd European Conference on Optical Communication*, Sep. 2016, pp. 1–3.

[8] L. Breyne, M. Verplaetse, C. Neumeyr, T. De Keulenaer, W. Soenen, X. Yin, *et al.*, "DSP-Free and Real-Time NRZ Transmission of 50 Gb/s Over 15-km SSMF and 64 Gb/s Back-to-Back With a 1.3-μm VCSEL", *Journal of Lightwave Technology*, vol. 37, no. 1, pp. 170–177, Jan. 2019.

[9] Y. Tamura, H. Sakuma, K. Morita, M. Suzuki, Y. Yamamoto, K. Shimada, *et al.*, "Lowest-Ever 0.1419-dB/km Loss Optical Fiber", in *Optical Fiber Communication Conference Postdeadline Papers*, Optical Society of America, 2017, Th5D.1. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=OFC-2017-Th5D.1.

[10] R. Giles, "Origins of the erbium-doped fiber amplifier", in *Frontiers in Optics 2016*, Optical Society of America, 2016, FTu1B.4. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=FiO-2016-FTu1B.4.

[11] R. J. Mears, L. Reekie, I. M. Jauncey, and D. N. Payne, "Low-Noise Erbium-doped Fibre Amplifier Operating at 1.54μm", *Electronics Letters*, vol. 23, no. 19, pp. 1026–1028, Sep. 1987.

[12] K. Iga, "Vertical-cavity surface-emitting laser (vcsel)", *Proceedings of the IEEE*, vol. 101, no. 10, pp. 2229–2233, Oct. 2013.

[13] F. Chang, *Datacenter Connectivity Technologies: Principles and Practice*, ser. River Publishers Series in Optics and Photonics. River Publishers, 2018. [Online]. Available: https://books.google.be/books?id=ooIstAEACAAJ.

[14] H. Hou, *Intel Silicon Photonics Update at Interconnect Day 2019*, Apr. 2019.

[15] F. Francini, D. Fontani, and P. Sansoni, "Transport of Light by Optical Fibers and Light Pipes", en, in *Sustainable Indoor Lighting*, ser. Green Energy and Technology, P. Sansoni, L. Mercatelli, and A. Farini, Eds., London: Springer, 2015, pp. 195–210.

[16] L. N. Binh, *Optical Fiber Communication Systems with MATLAB and Simulink Models*, 2nd Edition, ser. Optics and Photonics. Boca Raton, USA: CRC Press, 2015. [Online]. Available: https://doi.org/10.1201/b17781.

[17] L. Anet Neto, D. Erasme, N. Genay, P. Chanclou, Q. Deniel, F. Traore, *et al.*, "Simple Estimation of Fiber Dispersion and Laser Chirp Parameters Using the Downhill Simplex Fitting Algorithm", *Journal of Lightwave Technology*, vol. 31, no. 2, pp. 334–342, Jan. 2013.

[18] "ISO/IEC/IEEE International Standard - Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 3: Standard for Ethernet AMENDMENT 10: Media access control parameters, physical layers, and management parameters for 200 Gb/s and 400 Gb/s operation", *ISO/IEC/IEEE 8802-3:2017/Amd.10:2019(E)*, pp. 1–376, Feb. 2019.

[19] E. Säckinger, *Broadband Circuits for Optical Fiber Communication*. Hoboken, NJ, USA: John Wiley & Sons, Inc., Mar. 4, 2005. [Online]. Available: http://doi.wiley.com/10.1002/0471726400.

[20] W. Bogaerts, *Short Course 454 - OFC 2018: Hands-on: Introduction to Silicon Photonics Circuit Design*, Mar. 2018.

[21] M. Suzuki, Y. Noda, H. Tanaka, S. Akiba, Y. Kushiro, and H. Isshiki, "Monolithic integration of InGaAsP/InP distributed feedback laser and electroabsorption modulator by vapor phase epitaxy", *Journal of Lightwave Technology*, vol. 5, no. 9, pp. 1277–1285, Sep. 1987.

[22] imec, *iSiPP50G 300 mm Silicon Photonics Platform portfolio*, [Online; accessed May 2020]. [Online]. Available: https://www.imec-int.com/drupal/sites/default/files/2020-03/SILICON-PHOTONICS-V06.pdf.

[23] C. Doerr, L. Chen, D. Vermeulen, T. Nielsen, S. Azemati, S. Stulz, *et al.*, "Single-Chip Silicon Photonics 100-Gb/s Coherent Transceiver", in *Optical Fiber Communication Conference: Postdeadline Papers*, Optical Society of America, 2014, Th5C.1. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=OFC-2014-Th5C.1.

[24] M. C. Larson, J. Sonkoly, A. Mizrahi, M. Ayliffe, and J. Kemp, "InP vs Si Photonic Integrated Circuit Platforms for Coherent Data Center Interconnects", in *2018 European Conference on Optical Communication (ECOC)*, Sep. 2018, pp. 1–3.

[25] Z. Wang, A. Abbasi, U. Dave, A. De Groote, S. Kumari, B. Kunert, *et al.*, "Novel Light Source Integration Approaches for Silicon Photonics", *Laser & Photonics Reviews*, vol. 11, no. 4, p. 1 700 063, 2017. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/lpor.201700063.

[26] J. Zhang, G. Muliuk, J. Juvert, S. Kumari, J. Goyvaerts, B. Haq, *et al.*, "III-V-on-Si photonic integrated circuits realized using micro-transfer-printing", *APL Photonics*, vol. 4, no. 11, p. 110 803, 2019. [Online]. Available: https://doi.org/10.1063/1.5120004.

[27] G. Roelkens, W. Bogaerts, D. Taillaert, P. Dumon, L. Liu, S. Selvaraja, *et al.*, "Silicon nanophotonics: towards VLSI photonic integrated circuits", in *URSI (invited)*, 2008.

[28] C. R. Doerr, "Highly Integrated Monolithic Photonic Integrated Circuits", in *39th European Conference and Exhibition on Optical Communication (ECOC 2013)*, Sep. 2013, pp. 1–2.

[29] A. Abbasi, L. Abdollahi Shiramin, B. Moeneclaey, J. Verbist, X. Yin, J. Bauwelinck, *et al.*, "III–V-on-Silicon C-Band High-Speed Electro-Absorption-Modulated DFB Laser", *Journal of Lightwave Technology*, vol. 36, no. 2, pp. 252–257, Jan. 2018.

[30] R. Soref and B. Bennett, "Electrooptical effects in silicon", *IEEE Journal of Quantum Electronics*, vol. 23, no. 1, pp. 123–129, Jan. 1987.

[31] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson, "Silicon Optical Modulators", *Nature Photonics*, vol. 4, no. 8, pp. 518–526, Aug. 2010.

[32] M. Vanhoecke, A. Aimone, N. Argyris, S. Dris, R. Vaernewyck, K. Verheyen, *et al.*, "Segmented Optical Transmitter Comprising a CMOS Driver Array and an InP IQ-MZM for Advanced Modulation Formats", *Journal of Lightwave Technology*, vol. 35, no. 4, pp. 862–867, Feb. 2017.

[33] B. Milivojevic, S. Wiese, S. Anderson, T. Brenner, M. Webster, and B. Dama, "Demonstration of Optical Transmission at Bit Rates of up to 321.4Gb/s using Compact Silicon Based Modulator and Linear BiCMOS MZM Driver", *Journal of Lightwave Technology*, vol. 35, pp. 1–1, Jul. 2016.

[34] A. Aimone, I. G. Lopez, S. Alreesh, P. Rito, T. Brast, V. Höhns, *et al.*, "DAC-free ultra-low-power dual-polarization 64-QAM transmission with InP IQ segmented MZM module", in *2016 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2016, pp. 1–3.

[35] J. Sun, R. Kumar, M. Sakib, J. B. Driscoll, H. Jayatilleka, and H. Rong, "A 128 Gb/s PAM4 Silicon Microring Modulator With Integrated Thermo-Optic Resonance Tuning", *Journal of Lightwave Technology*, vol. 37, no. 1, pp. 110–115, Jan. 2019.

[36] Y. Ban, J. Verbist, M. Vanhoecke, J. Bauwelinck, P. Verheyen, S. Lardenois, *et al.*, "Low-Voltage 60Gb/s NRZ and 100Gb/s PAM4 O-Band Silicon Ring Modulator", in *2019 IEEE Optical Interconnects Conference (OI)*, Apr. 2019, pp. 1–2.

[37] S. Gupta, S. A. Srinivasan, M. Pantouvaki, H. Chen, P. Verheyen, G. Lepage, *et al.*, "50GHz Ge Waveguide Electro-Absorption Modulator Integrated in a 220nm SOI Photonics Platform", in *Optical Fiber Communication Conference*, Optical Society of America, 2015, Tu2A.4. [Online]. Available: `http://www.osapublishing.org/abstract.cfm?URI=OFC-2015-Tu2A.4`.

[38] S. A. Srinivasan, P. Verheyen, R. Loo, I. D. Wolf, M. Pantouvaki, G. Lepage, *et al.*, "50Gb/s C-band GeSi Waveguide Electro-Absorption Modulator", in *Optical Fiber Communication Conference*, Optical Society of America, 2016, Tu3D.7. [Online]. Available: `http://www.osapublishing.org/abstract.cfm?URI=OFC-2016-Tu3D.7`.

[39] S. A. Srinivasan, C. Porret, E. Vissers, P. Geiregat, D. V. Thourhout, R. Loo, *et al.*, "High-contrast quantum-confined stark effect in ge/sige quantum well stacks on si with ultra-thin buffer layers", in *CLEO Pacific Rim Conference 2018*, Optical Society of America, 2018, Th3C.1. [Online]. Available: `http://www.osapublishing.org/abstract.cfm?URI=CLEOPR-2018-Th3C.1`.

[40] Clement Porret and Srinivasan Ashwyn Srinivasan and Sadhishkumar Balakrishnan and Peter Verheyen and Paola Favia and Hugo Bender and Patrick Ong and Roger Loo and Joris Van Campenhout and Marianna Pantouvaki, "O-band GeSi quantum-confined Stark effect electro-absorption modulator integrated in a 220nm silicon photonics platform", in *VLSI Technology 2020*, 2020.

[41] A. Abbasi, B. Moeneclaey, J. Verbist, X. Yin, J. Bauwelinck, G. Duan, *et al.*, "Direct and Electroabsorption Modulation of a III–V-on-Silicon DFB Laser at 56 Gb/s", *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 23, no. 6, pp. 1–7, Nov. 2017.

[42] T. Chan and W. Way, "Pam4 modulation using electro-absorption modulated lasers", in *Datacenter Connectivity Technologies: Principles and Practice*, F. Chang, Ed., River Publishers, 2018, ch. 4, pp. 173–212.

[43] H. Ramon, M. Vanhoecke, J. Verbist, W. Soenen, P. De Heyn, Y. Ban, *et al.*, "Low-Power 56Gb/s NRZ Microring Modulator Driver in 28nm FDSOI CMOS", *IEEE Photonics Technology Letters*, vol. 30, no. 5, pp. 467–470, Mar. 2018.

[44] H. Ramon, J. Lambrecht, J. Verbist, M. Vanhoecke, S. A. Srinivasan, P. De Heyn, *et al.*, "70 Gb/s Low-Power DC-Coupled NRZ Differential Electro-Absorption Modulator Driver in 55 nm SiGe BiCMOS", *Journal of Lightwave Technology*, vol. 37, no. 5, pp. 1504–1514, Mar. 2019.

[45] H. Li, G. Balamurugan, M. Sakib, J. Sun, J. Driscoll, R. Kumar, *et al.*, "A 112 Gb/s PAM4 Silicon Photonics Transmitter With Microring Modulator and CMOS Driver", *Journal of Lightwave Technology*, vol. 38, no. 1, pp. 131–138, Jan. 2020.

[46] A. Samani, E. El-Fiky, M. Morsy-Osman, R. Li, D. Patel, T. Hoang, *et al.*, "Silicon Photonic Mach–Zehnder Modulator Architectures for on Chip PAM-4 Signal Generation", *Journal of Lightwave Technology*, vol. 37, no. 13, pp. 2989–2999, Jul. 2019.

[47] H. Li, G. Balamurugan, M. Sakib, J. Sun, J. Driscoll, R. Kumar, *et al.*, "A 112 gb/s pam4 silicon photonics transmitter with microring modulator and cmos driver", *J. Lightwave Technol.*, vol. 38, no. 1, pp. 131–138, Jan. 2020. [Online]. Available: `http://jlt.osa.org/abstract.cfm?URI=jlt-38-1-131`.

[48] J. Verbist, J. Lambrecht, M. Verplaetse, J. Van Kerrebrouck, A. Srinivasan, P. De Heyn, *et al.*, "DAC-Less and DSP-Free 112 Gb/s PAM-4 Transmitter Using Two Parallel Electroabsorption Modulators", *Journal of Lightwave Technology*, vol. 36, no. 5, pp. 1281–1286, Mar. 2018.

[49] J. Verbist, M. Verplaetse, S. A. Srivinasan, P. De Heyn, T. De Keulenaer, R. Pierco, *et al.*, "First Real-Time 100-Gb/s NRZ-OOK Transmission over 2 km with a Silicon Photonic Electro-Absorption Modulator", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[50] H. Chen, P. Verheyen, P. D. Heyn, G. Lepage, J. D. Coster, S. Balakrishnan, *et al.*, ''-1 V bias 67 GHz bandwidth Si-contacted germanium waveguide p-i-n photodetector for optical links at 56 Gbps and beyond", *Opt. Express*, vol. 24, no. 5, pp. 4622–4631, Mar. 2016. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-24-5-4622.

[51] J. Verbist, J. Lambrecht, B. Moeneclaey, J. Van Campenhout, X. Yin, J. Bauwelinck, *et al.*, "40-Gb/s PAM-4 Transmission Over a 40 km Amplifier-Less Link Using a Sub-5V Ge APD", *IEEE Photonics Technology Letters*, vol. 29, no. 24, pp. 2238–2241, Dec. 2017.

[52] H. T. Chen, J. Verbist, P. Verheyen, P. De Heyn, G. Lepage, J. De Coster, *et al.*, "25-Gb/s 1310-nm Optical Receiver Based on a Sub-5-V Waveguide-Coupled Germanium Avalanche Photodiode", *IEEE Photonics Journal*, vol. 7, no. 4, pp. 1–9, Aug. 2015.

[53] S. A. Srinivasan, J. Lambrecht, M. Berciano, S. Lardenois, P. Absil, J. Bauwelinck, *et al.*, "Highly Sensitive 56 Gbps NRZ O-band BiCMOS-Silicon Photonics Receiver using a Ge/Si Avalanche Photodiode", in *Optical Fiber Communication Conference (OFC) 2020*, Optical Society of America, 2020, W4G.7. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=OFC-2020-W4G.7.

[54] V. NWS, *Microsoft en Facebook gaan onderzeese kabel aanleggen tussen VS en Europa*, 2016 (accessed: 08.02.2020). [Online]. Available: https://www.vrt.be/vrtnws/nl/2016/05/27/microsoft_en_facebookgaanonderzeesekabelaanleggentussenvsen1-2666843/.

[55] Samtec, *Twinax flyover system*, accessed: 08.02.2020. [Online]. Available: https://www.samtec.com/s2s/system-optimization/twinax-flyovers.

[56] Molex, *Bipass*, accessed: 08.02.2020. [Online]. Available: https://www.molex.com/molex/products/family/bipass_io_and_backplane_cable_assemblies.

[57] "Iso/iec/ieee international standard - information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements - part 3: Standard for ethernet amendment 10: Media access control parameters, physical layers, and management parameters for 200 gb/s and 400 gb/s operation", *ISO/IEC/IEEE 8802-3:2017/Amd.10:2019(E)*, pp. 1–376, Feb. 2019.

[58] B. Moeneclaey, J. Verbrugghe, J. Lambrecht, E. Mentovich, P. Bakopoulos, J. Bauwelinck, *et al.*, "Design and experimental verification of a transimpedance amplifier for 64-gb/s pam-4 optical links", *Journal of Lightwave Technology*, vol. 36, no. 2, pp. 195–203, Jan. 2018.

[59] S. Gudyriev, C. Kress, H. Zwickel, J. N. Kemal, S. Lischke, L. Zimmermann, *et al.*, "Coherent epic receiver for 64 gbaud qpsk in 0.25 Photonic BiCMOS Technology", *Journal of Lightwave Technology*, vol. 37, no. 1, pp. 103–109, Jan. 2019.

[60] C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, *et al.*, "Single-chip microprocessor that communicates directly using light", *Nature*, vol. 528, no. 7583, pp. 534–538, Dec. 2015.

[61] S. L. (Intel), *Ayar Labs and Intel demo FPGA with Optical Transceivers in Darpa PIPES Project: 2Tbps now, >100 Tbps is the Goal.* [Online; accessed April 2020], Mar. 2020. [Online]. Available: https://blogs.intel.com/psg/ayar-labs-and-intel-demo-fpga-with-optical-transceivers-in-darpa-pipes-project-2-tbps-now-100-tbps-is-the-goal/.

[62] IHP, *Sg25h5*$_e pic : Amonolithicphotonicsbicmostechnology$, [Online; accessed April 2020]. [Online]. Available: %7Bhttps://www.ihp-microelectronics.com/en/departments/technology/si-photonics.html%7D.

[63] GLOBALFOUNDRIES, *90wg and 45clo: Gf's 90.1667emnm and 45.1667emnm rf soi cmos nodes with siph*, [Online; accessed April 2020]. [Online]. Available: %7Bhttps://www.globalfoundries.com/sites/default/files/product-briefs/harness_the_power_of_light_mar2020b.pdf%7D.

[64] P. Moorhead, *GlobalFoundries Has Quietly Become A Player In Silicon Photonics Manufacturing*, [Online; accessed April 2020], Mar. 2020. [Online]. Available: https://www.forbes.com/sites/

moorinsights / 2020 / 03 / 31 / globalfoundries - has - quietly - become - a - player - in - silicon - photonics - manufacturing/#20c1776a8bdb.

3

Real-Time 100 Gb/s NRZ and EDB Transmission with a GeSi Electro-Absorption Modulator for Short-Reach Optical Interconnects

In this chapter, we demonstrate the first silicon-based EAM, in combination with an in-house developed SiGe BiCMOS transceiver chipset, capable of transmitting single-lane 100 Gb/s non-return-to-zero in real-time. Transmission over 500 m of standard single mode fiber and 2 km of non-zero dispersion shifted fiber is demonstrated, assuming a forward-error coding scheme is used with a bit-error ratio limit of 3.8×10^{-3}. Due to the high line rate, transmission over longer fiber spans was limited by the chromatic distortion in the fiber. As a possible solution, electrical duobinary modulation is proposed as it is more resilient to this type of fiber distortion by reducing the required optical bandwidth. We show improved performance for longer fiber spans with a 100 Gb/s electrical duobinary link, resulting in real-time sub-FEC operation over more than 2 km of standard single-mode fiber without any digital signal processing. Finally, the possibility of a 100 Gb/s EAM-to-EAM link is investigated.

This chapter is based on the journal paper with the same title that was

published in Journal of Lightwave Technology (Dec. 2017) [1]. It was an invited extension of our work presented as post deadline paper at OFC 2017 [2]

3.1 Introduction

In order to meet the growing bandwidth requirements, data centers will soon require short-reach optical interconnects to operate at 100 Gb/s and beyond. Recently, this has led to an evolution from 100 Gb/s Ethernet to 400 Gb/s Ethernet (400 GbE). At the time of this research several possible implementations were still under discussion in the standardization committees [3].

Among the different approaches, the 4×100 Gb/s configuration -either through coarse wavelength division multiplexing or multiple fibers- surfaced as the most viable solution. A four lane 100 Gb/s non-return-to-zero (NRZ) scheme could provide an elegant solution towards a compact 400 GbE transceiver, allowing a more compact and low-cost transceiver through lower lane counts, while maintaining the low complexity of on-off-keying-based electronics. While the standards now focus on 50 Gbaud PAM-4 solutions, the 100G NRZ approach could still be used in the context of active optical cables or serve as a baseline for 100 Gbaud PAM-4 for 200G/lambda transceivers at 800GbE and 1.6TbE.

Previously, several 100 Gb/s single-lane transmissions have been realized using four level pulse amplitude modulation (PAM-4)[4–7], discrete multitone (DMT)[8, 9] and electrical 3-level duobinary (EDB)[10–12]. However, many of these experiments, especially for PAM-4 and DMT, still rely on complex digital signal processing (DSP) at the RX and/or TX-side, typically done offline. Silicon photonics would be ideally suited to implement such a scheme as it can provide compact and low-cost transceivers, although scaling to 100 Gb/s lane rates has proven to be difficult for silicon based-transceivers, with currently only non-real-time demonstrations which still have to rely on large traveling wave structures and/or extensive DSP [13, 14].

Nevertheless, some examples of true real-time 100 Gb/s serial rates without DSP have been demonstrated recently on other platforms. In [7], a discrete Mach-Zehnder modulator was operated at 100 Gb/s with custom designed TX and RX consuming 8.6 W. Recently, a real-time 100 Gb/s EDB modulation was reported in [11], where an InP-based traveling-wave EAM with integrated distributed feedback laser was used to transmit below the hard-decision forward error coding limit (HD-FEC) of 3.8×10^{-3}.

The transmission line design of the electrode does not only increase the overall device size when compared to a lumped driven modulator, but also necessitates a power consuming 50 Ω termination. The same transmitter as in [11] was used for a real-time 100 Gb/s NRZ link in [12]. Unfortunately, the transceiver modules were developed for metro networks, leading to unrealistic form factors and power consumptions for use in short-reach optical interconnects. Finally, an impressive BER down to 10^{-10} with 100 Gb/s NRZ on a polymer Mach-Zehnder modulator was presented in [15], again with traveling wave electrodes and 50 Ω termination. However, the proposed transceiver does pose several drawback in terms of cost, power and footprint, when envisioned as a device for short-reach optical interconnects. Two co-packaged InP-based electrical multiplexers are required, offering a 6 V_{pp} differential voltage swing to drive the 1.1 cm long modulator. This results in a total power consumption -excluding the laser- of 6.85 W for the 6.5×2 cm transmitter module. The proposed 4×100 PIN-DEMUX receiver adds another 5.5 W and measures 4×6.9 cm. Even with the addition of transimpedance amplifier (TIA), removal of the erbium-doped optical amplifier (EDFA) in the RX from the link might be difficult as the the maximal transmitted output power is limited to -9.5 dBm.

In this chapter, we present a real-time, single-lane and serial 100 Gb/s NRZ-OOK link with a silicon-based electro-absorption modulator (EAM) in combination with an in-house developed transmitter (TX-IC) and receiver (RX-IC) chipset in a SiGe BiCMOS technology. The extremely compact GeSi EAM was fabricated on a 200 mm silicon-on-insulator platform without any traveling-wave electrodes and/or power-dissipating terminations. Transmission of 100 Gb/s NRZ over 500 m of standard single-mode fiber (SSMF) and 2 km of non-zero dispersion shifted fiber (NZ-DSF) is reported. We also investigate the performance of EDB modulation in the same link. Successful real-time transmission at 100 Gb/s EDB, assuming a HD-FEC, is demonstrated over more than 2 km of SSMF. These are the first real-time chip-to-chip demonstrations of a 100 Gb/s NRZ or EDB link with a silicon-based waveguide modulator without the need for temperature control, material post-processing or complex DSP.

3.2 Components for 100 Gb/s short-reach optical interconnects

At bitrates of 100 Gb/s and higher, careful design of the both electrical and optical components is needed, especially when envisioning a limited

power-budget and form factor. In section 3.2.1, the electrical transmitter and receiver which provide the capability of equalizing and decoding 100 Gb/s NRZ or EBD in real-time are discussed. Next, the design, characterization and operation of the silicon-based EAM is presented in section 3.2.2.

3.2.1 Electrical Transceiver

To generate and receive 100 Gb/s data in real-time an in-house developed transmitter and receiver were used. A first generation transmitter, fabricated in a 130 nm SiGe BiCMOS technology was used for the NRZ experiments (at 1601.5 nm). For the EDB measurements (at 1560 nm), this transmitter was replaced by a new version of the IC implemented in a 55 nm SiGe BiCMOS process with improved bandwidth and power-consumption. The transmitter IC (TX IC) consists of 2 main building blocks: a 4-to-1 multiplexer (MUX) which generates a 100 Gb/s data stream out of four 25 Gb/s streams and a six-tap analog feedforward equalizer (FFE), as can be seen in Fig.3.1. The equalizer was implemented at the TX-side to reduce the dynamic range requirements on the RX-IC, at the cost of necessitating a linear output buffer after the FFE. An other possible benefit is the exclusion of noise-shaping by an FFE at the RX-side. The main drawback is the automatic optimization of the FFE in a practical system. This would require some form of back-channel (albeit at much lower speeds) to update the FFE settings, possibly from a least-mean-square engine located at the RX [16]. Fig.3.2 demonstrates the effect of the FFE when set for a 100 Gb/s NRZ transmission over a short coaxial RF-cable (Fig.3.2a), predistorted for 100 Gb/s optical back-to-back (B2B) NRZ transmission (Fig. 3.2b) and the resulting optical eye captured with a high-speed photodiode (Fig. 3.2c). The FFE taps are symbol-spaced at 9-10 ps allowing us to equalize up to 50 GHz, but only over 60 ps. At a serial rate of 100 Gb/s, the 130 nm TX-IC consumes 1 W and the 55 nm version 0.75 W. The dies measure 1.5×4.5 mm and 1×3.8 mm, respectively. In both cases the MUX and the FFE + output driver split the total power consumption in a 35%-65% way.

To decode the received signal, the receiver IC (RX-IC) presented in [17] was used. The chip was fabricated in a 130 nm SiGe BiCMOS process and performs two main tasks: it samples and decodes the incoming signal and it demultiplexes the full rate data stream into four quarter rate streams as shown in Fig. 3.1. Because the RX-IC was primarily designed for the reception of duobinary signals, there are two independent parallel comparators followed by two limiting amplifiers (LA) to sample the upper

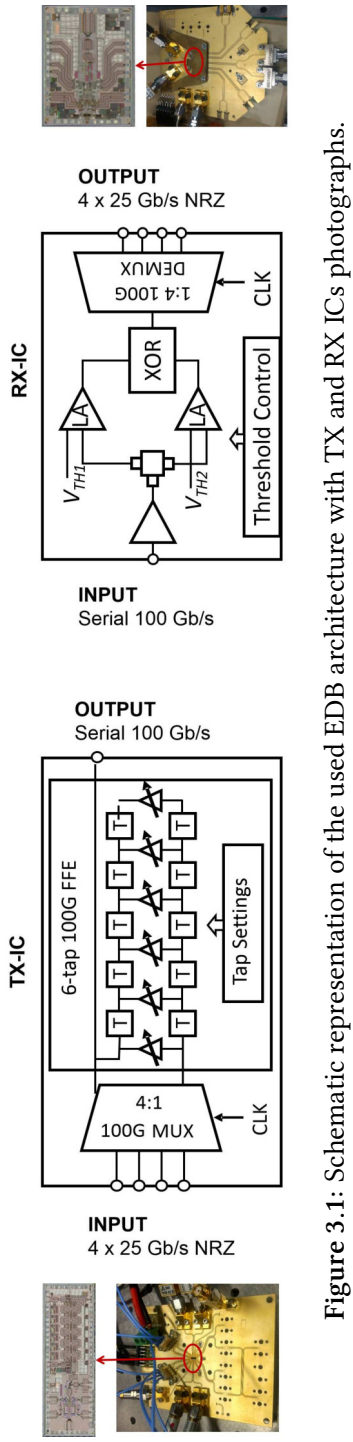

Figure 3.1: Schematic representation of the used EDB architecture with TX and RX ICs photographs.

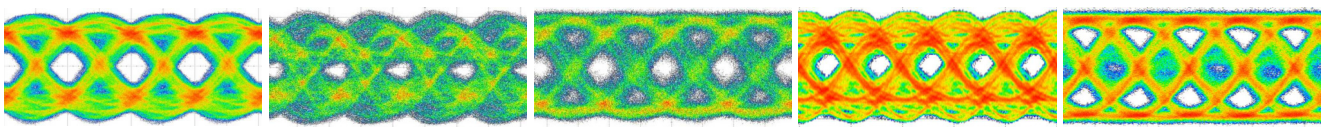

Figure 3.2: (a) at output of transmitter IC and optimized for electrical NRZ transmission,(b) predistorted for optical NRZ transmission, (c) the resulting optical NRZ eye-diagram after PIN-PD, (d) predistorted for EDB transmission and (e) resulting optical EDB eye-diagram after PIN-PD

and the lower eye of the typical 3-level duobinary eye. Next, a XOR-port is used to decode and convert the streams from the sampled upper and lower eye data back into the original pre-coded NRZ format. Of course, if one of the comparator thresholds is fixed HIGH the XOR-port becomes functionally transparent and the receiver reduces to a conventional NRZ decoder. This allows us to transmit and receive duobinary and NRZ signals in real-time with the same transceivers. In section 3.4 we will briefly discuss why it might be interesting to switch from EDB to NRZ depending on the optical link. The chip measures 2×2.6 mm and consumes less than 1.2 W, of which the DEMUX contributes 0.7 W, at a serial rate of 100 Gb/s. In this version, no clock-and-data-recovery circuit is available so the alignment of the sampling clock with the optimal sample time was done manually with an external tunable time delay.

The overall bandwidth of the transceiver chipset is dominated by the bandwidth of the input amplifier of the RX-IC at 41 GHz. This suffices for duobinary modulation, but requires quite some high-pass shaping by the FFE for NRZ links. Nevertheless, error-free operation over a short short electrical link was obtained for both modulation formats, when connecting the transmitter and the receiver IC with RF coax-cable. A continuous BER measurement revealed a BER of 1×10^{-12} for NRZ modulation and 1×10^{-13} for EDB modulation. At 100 Gb/s, the transceiver chipset is able to serialize, equalize, decode and deserialize for a combined electrical power consumption of 1.95 W (when using the 55 nm TX-IC). This amounts to an energy/bit of 19.5 pJ/bit, excluding the power consumption of the RF amplifier needed to drive the EAM with $2\,V_{pp}$ (i.e. an additional 24 pJ/bit). Currently, an external RF-Amplifier is needed as the TX-IC was designed to transmit over electrical backplanes, requiring much less voltage swing. However, adding 1 or 2 additional amplifying stages in the TX-IC design should allow to increase this to approximately $2\,V_{pp}$, which is quite feasible in the used BiCMOS technology with a nominal supply voltage of 2.5 V. Although this would inevitably result in a significant increase in the total power consumption of the transceiver (estimated to be an additional 200 mW or 2 pJ/bit), it would still be much less than the 24 pJ/bit used by the external RF amplifier.

3.2.2 GeSi Electro-Absorption Modulator

The high-speed waveguide electro-absorption modulator was fabricated in imec's silicon photonics platform on a 200 mm silicon-on-insulator wafer with 220 nm top Si thickness and consists of a 600 nm wide and 80 μm long germanium waveguide with embedded p-i-n-junction. The cross-section

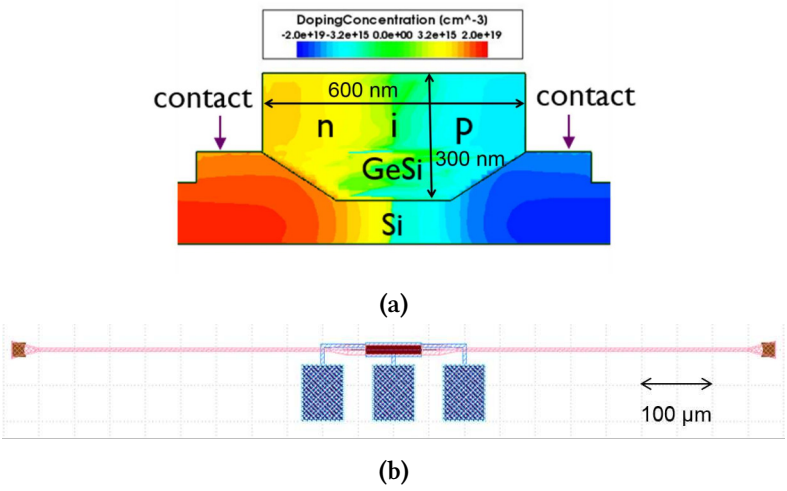

Figure 3.3: (a) Cross-section of the GeSi waveguide EAM with indication of doping concentrations; (b) Layout for fabrication of the proposed 80 µm long EAM terminated by two fiber-to-chip grating couplers.

of the GeSi EAM is shown in Fig.3.3. Modulation is based on the Franz-Keldysh effect, where the effective bandgap of the GeSi shifts when an electrical field is present[18]. Incorporating ∼0.8% of Si shifts the band edge sufficiently to allow operation around 1550 nm compared to a pure Ge EAM operating around 1610 nm [19]. More information regarding the design and fabrication of a 40 µm long version of this EAM can be found in [18]. Light is coupled in and out of the waveguide structure through fiber-to-chip grating couplers with an insertion loss of ∼6 dB per coupler. The EAM was operated around 1560 nm for EDB experiments and around 1600 nm for both NRZ and EDB experiments. At a 2 V_{pp} swing and a bias of -2 V, the GeSi EAM has a junction capacitance of ∼15 fF, leading to a dynamic average energy per bit of less than 15 fJ/bit. For a fair comparison, the static power consumption of the EAM should also be taken into consideration. For an in-waveguide power of 6 dBm at a comparable bias of -2.05 V, the EAM produced a DC photocurrent of approximately 3.8 mA, resulting in a static average energy per bit of 76 fJ/bit. Combined, this amounts to a total average energy/bit of less than 91 fJ/bit of the EAM during all real-time NRZ experiments at 1601.5 nm. During the C-band experiments, the modulator generated a DC photocurrent of 2.39 mA at a bias of -0.85 V, reducing the static energy per bit to 20 fJ/bit. In both cases, the power consumption in the transmitter is dominated by that of

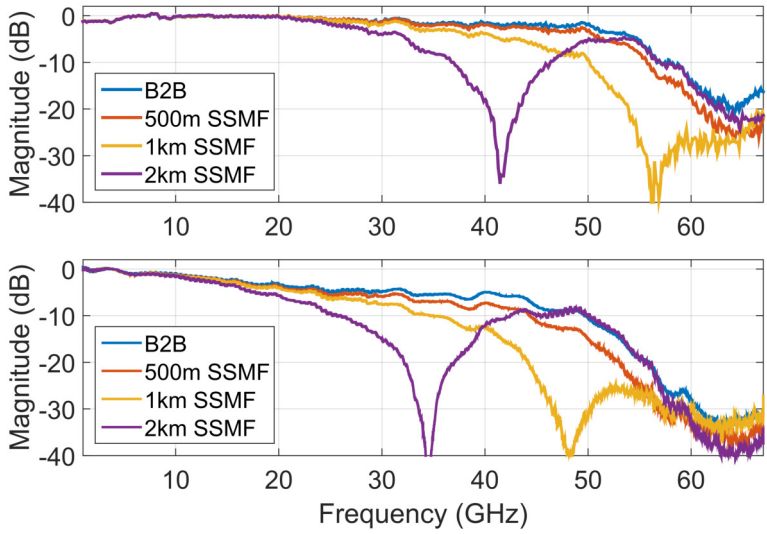

Figure 3.4: Small-signal frequency response of the optical link consisting of the RF amplifier, GeSi EAM and a 50 GHz commercial PD for fiber spans up to 2 km at 1560 nm (top) and 1600 nm (bottom).

the electrical front-end required to drive it.

3.2.3 Chromatic Distortion in the Fiber Channel at 100 Gb/s

Not only the electrical transceiver chipset and the optical modulator are important parts of a 100 Gb/s link, the fiber channel itself plays a significant role when operating at wavelengths in C- and L-band. At those wavelengths the relatively large chromatic dispersion coefficient manifests itself as notches in the frequency response of the optical link, limiting the overall bandwidth. The small-signal frequency response of the modulator driven by a 50 GHz RF amplifier and received by a 50 GHz p-i-n diode is given in Fig. 3.4 for different fiber spans (0, 500 m, 1 km and 2 km) at 1560 nm and 1600 nm. At 1560 nm 2 km of SSMF introduces a notch around 41 GHz which degrades the frequency response in area of approximately ± 14 GHz around this notch. As expected, at 1600 nm the notches are located at even lower frequencies due to the steadily increasing dispersion coefficient from C-band to L-band. This poses severe limitations on the maximal fiber span at 100 Gb/s without resorting to chromatic distortion (CD) compensation techniques such as dispersion shifted or compensated fiber links.

3.3 Experiment Setup

The experiment setup is illustrated in Fig.3.5. A Xilinx VirteX FPGA board generates four $2^7 - 1$ long pseudo-random bit streams (PRBS) at 25 Gb/s, which are serialized to a 100 Gb/s single line rate with required delays to form a $2^7 - 1$ long stream at 100 Gb/s. In these first experiments, the possible penalty on the link performance with longer PRBS sequence was not yet investigated. Next, the six-tap analog equalizer in the TX-IC is set to compensate the frequency roll-off and other non-idealities of the following components in the link. Even though the tap settings were optimized for each experiment, a configuration with one pre-cursor, one main and 4 post-cursor taps was found to give good all-round performance and was kept for all subsequent experiments. A 50 GHz RF-amplifier with internal bias-T at the output is used to apply the pre-emphasized signal from the TX-IC with a 2 V_{pp} swing via an RF-probe to the bondpads of the EAM. As the EAM is a small capacitive load driven by a 50 Ω, we expect the effective voltage over the modulator to double (especially at lower frequencies) when compared to driving a 50 Ω load. However, we can use the FFE to shape the data with the inverse characteristic (i.e. a high-pass filter) to compensate the combined effect of frequency dependent losses and reflections (∼-6 dB at 50 GHz for the channel after the TX-IC output until the EAM). The measured peak-to-peak voltage in a 50 Ω load (a 70 GHz sampling oscilloscope) of the TX-IC output was approximately 200 mV_{pp}, when set to drive the EAM (Fig. 3.2-b). With 20 dB gain form the RF amplifier, we arrive at a swing of 2 V_{pp}.

During the NRZ measurements, light at 1601.5 nm is sent into the EAM with an in-waveguide power around 6 dBm, while 2 dBm power at 1560 nm was used for the EDB experiments. The EAM was biased at -1.85 V for back-to-back L-band NRZ links and slightly higher at -2.05 V during transmission experiments, resulting in a photocurrent of roughly 3.6 mA and 3.8 mA, respectively. For EDB modulation in C-band the bias was set to -0.65 V for B2B links and again increased slightly for optimal performance to -0.85 V during transmission experiments. With these settings we measured a dynamic extinction ratio of ∼6 dB at 1601.5 nm and a bit more than 7 dB at 1560 nm. The insertion loss for both modes of operation was estimated around ∼6 dB. During all experiments, the EAM was operated at room temperature without any temperature control. A commercial 50 GHz IIIV-based p-i-n photodiode (PD) converted the optical signal back into the electrical domain. As no transimpedance amplifier (TIA) with sufficient bandwidth (i.e.>50 GHz) was available, an erbium-doped-amplifier (EDFA) was used boost the maximal input power to the

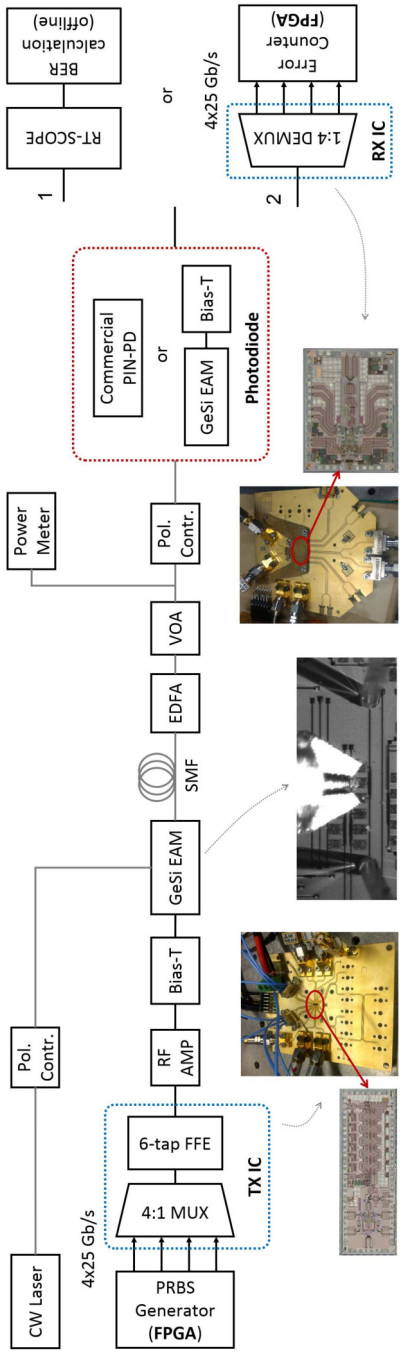

Figure 3.5: Experiment setup of real-time 100 Gb/s NRZ/EDB optical link.

PD. This was needed as the sensitivity of the RX-IC is 18 mV_{pp} for a BER of 1×10^{-12} (for EDB transmissions). The EDFA could be removed from the link with the addition of a TIA and/or by replacing the fiber-to-chip grating couplers with low-loss edge-couplers (<2 dB/coupler), as will be discussed in section 3.4.1. Finally, the received bitstream is decoded for respectively NRZ or EDB by setting the right comparator levels as discussed in section 3.2.1, and deserialized into four 25 Gb/s NRZ streams and fed back to the FPGA for real-time error detection.

In section 3.4.3, the commercial PD was replaced by an identical copy of the GeSi on a second die, acting as a photodiode. These experiments, as well as the reference curves in section 3.4.2, were done by capturing the signal from the photodetector (commercial PD or second EAM) by a real-time 160 GSa/s oscilloscope, after which the BER was calculated offline.

3.4 Results and Discussion

3.4.1 100 Gb/s NRZ Transmission

In a first experiment, real-time NRZ transmission at 1601.5 nm was carried out using the electrical transceiver discussed in section 3.2.1 and a commercial PD as an optical receiver as shown in Fig. 3.5. The real-time BER curves for transmission over several fiber spans can be seen in Fig. 3.6a and examples of received NRZ eyes captured by a 70 GHz sampling oscilloscope are shown in 3.2. As the RX-IC poses the main bandwidth limitation in the link and provides only deserialized quarter-rate outputs, setting the FFE is done in a two-step approach. First, the received eye is optimized through visual inspection on a sampling oscilloscope. Next, the resulting tap settings are used as a starting point for further manual optimization by minimizing the BER of the quarter-rate outputs. For a B2B link, a BER of below 6×10^{-9} was obtained at an average optical power of 8.3 dBm in the PD. The hard-decision forward error coding limit (HD-FEC: 3.8×10^{-3} for 7% overhead) was reached for an average power above -0.6 dBm. Although other FEC standards exist (e.g. KR4: 5.2×10^{-5} and KP4: 2×10^{-4} for 100G PAM-4 intra-data center interconnects), this FEC limit was chosen to provide quick and easy comparisons between the different experiments in this chapter, as well as in literature, as currently no standards exist for 100 Gb/s NRZ links. As shown in Fig. 3.4, the chromatic distortion at around 1600 nm severely degrades of the frequency response and reduces the overall bandwidth of the link. Nevertheless, we still manage to obtain a BER below 2×10^{-5} for 500 m of SSMF. Sub-FEC operation is realized for >1.5 dBm, resulting

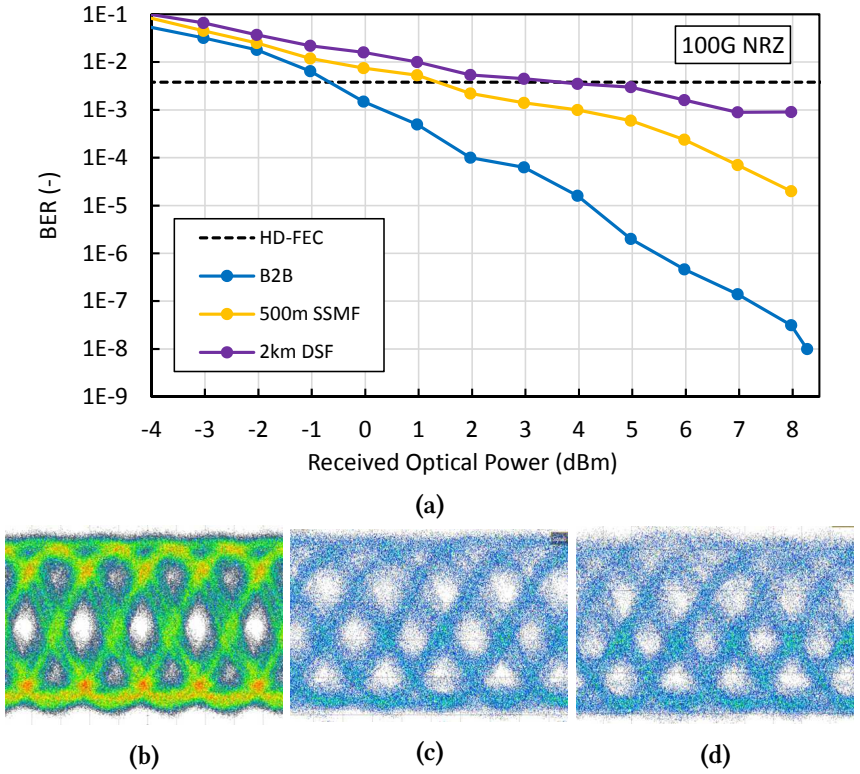

Figure 3.6: (a) Real-time BER curves and received eye diagrams for 100 Gb/s NRZ for (b) B2B, (c) 500 m of SSMF and (d) 2 km of DSF ($\sim$ 8 ps/nm.km) at 1601.5 nm.

in a power penalty of 2.1 dB compared to B2B. Lastly, transmission over 2 km of non-zero dispersion-shifted fiber (NZ-DSF; dispersion coefficient of $\sim$8 ps/(nm.km)) is investigated, requiring 3 dBm of optical power to drop below the HD-FEC limit and saturating in a BER just below 1×10^{-3} for higher powers. For B2B and 500 m SSMF, no saturation of the BER has yet occurred. Because the total dispersion of the 2 km NZ-DSF at 1601.5 nm is approximately equal to that of 1 km of standard SMF $\sim$16 ps/nm), transmission over 1 km should result in comparable BERs, but was not measured. As the average in-fiber power after the modulator due to the insertion loss of the EAM and two grating couplers was around -5 dBm, we would only need to improve our link budget with 4.5 dB to reach the HD-FEC limit for a B2B transmission and 6.5 dB for a 500 m transmission. One possible solution would be to replace the fiber-to-chip

grating couplers (∼6 dB/coupler) with low-loss edge-couplers (typically < 2 dB/coupler)[20]. This would boost the power budget by almost 10 dB, allowing us to remove the EDFA from the setup and realizing an amplifier-less link, assuming the increased input-power does not significantly change the behavior of the EAM. We would also like to indicate that current system experiments did not include a transimpedance amplifier. Adding a TIA would give a substantial improvement in sensitivity, which should be sufficient to stay under HD-FEC for an average output power of 0 dBm. This could be directly obtained by replacing the output grating coupler with a low-loss edge coupler, without changing the optical input power to the EAM.

3.4.2 100 Gb/s Duobinary Transmission

In order to realize successful transmission up to 2 km of SSMF, a couple of changes are made to the experiment setup. First, the FFE is re-optimized to shape the transmitted data into an electrical duobinary format, as illustrated in Fig. 3.2d and Fig. 3.2e. Next, to further minimize the effect of the CD, the operational wavelength is shifted to C-band (1560 nm). A 3 dB attenuator is added to the output of the RF-amplifier to shield this 50 Ω driver from the reflections due to the mainly capacitive loading by the EAM, especially at lower frequencies. As consequence, the FFE can be used more efficiently, as it has to put less effort in attenuating these frequency components when shaping the channel. Lastly, as discussed in section 3.2.1 a newer and faster, but functionally identical version of the TX-IC was used during these experiments.

In a first experiment as a reference, BER curves for 100 Gb/s EDB (shown in Fig. 3.7a) were measured for 0, 1 and 2 km of SSMF by capturing >10 million symbols with a 160 GSa/s real-time sampling oscilloscope and calculating the BER offline. The optimal thresholds were determined via a histogram over a thousand symbols and was swept over the possible sampling times after interpolation of the received data with a factor 10 (i.e. 16 samples/symbol at 100 Gb/s). The data is aligned and compared to the transmitted signal. No other digital signal processing or filtering was used.

Even though the eyes after 2 km have slightly degraded compared to 0 km and 1 km, we still have decently open eyes, as can be seen in Fig. 3.7a and operation down to a BER of 1×10^{-5} is possible for all fiber spans up to 2 km. Sub-FEC operation is obtained for average optical powers above 0.6 dBm for all lengths of fiber. No clear saturation of the BER is apparent

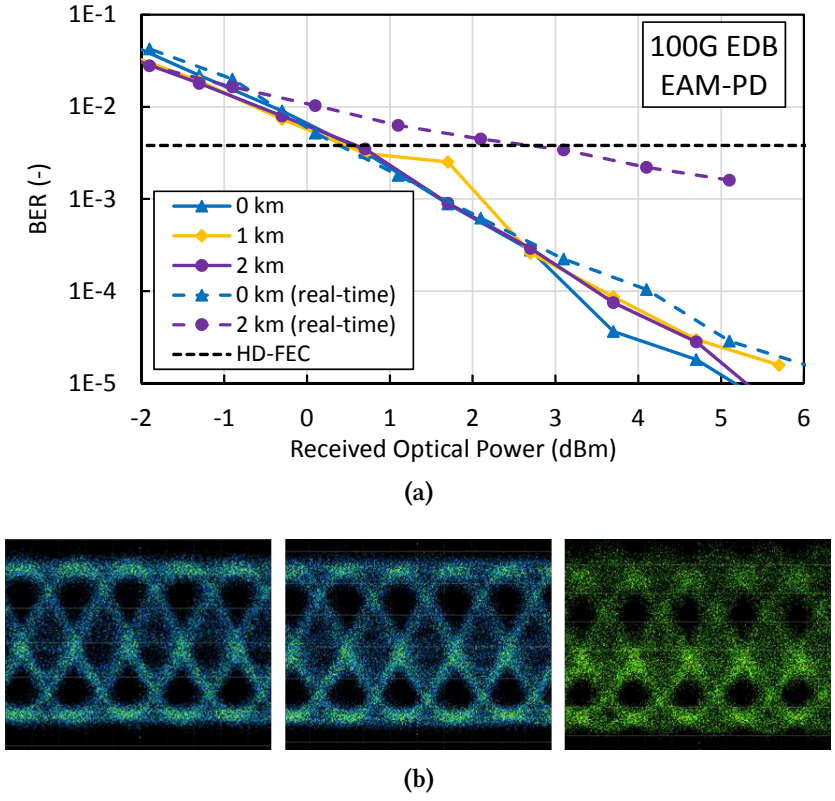

Figure 3.7: (a) Measured BER curves for duobinary modulation at 1560 nm. The full lines (-) correspond to offline calculated BERs from data captured with a real-time oscilloscope and the dotted lines (–) are real-time end-to-end measurements with the electrical receiver. (b) Examples of a 100 Gb/s EDB eye diagrams at 5 dBm of average optical power after 0, 1 and 2 km of SSMF.

yet.

In a second experiment, real-time transmission was again investigated. For a B2B link, the BER curve is fairly comparable to that of the offline measured BER curve up to 3 dBm, after which a penalty of $\sim$ 1 dB appears for higher powers. With 2 km of SSMF the penalty with respect to the reference curves is much larger ($\sim$2.1 dB at HD-FEC) and we can see the onset of BER saturation emerging. Nevertheless, we still manage to obtain successful sub-FEC operation up to 2 km of SSMF, a clear improvement

compared to NRZ modulation discussed in section 3.4.1.

With the longest typical fiber distances in hyperscale datacenter limited to 2 km, an EDB modulation based transceiver would be ideally suited for this type of interconnect, where the increased complexity of transitioning from NRZ-based transceivers to an EDB-based transceivers is warranted to cover these distances without having to resort to more complex schemes (e.g. PAM-4) or DSP. However, in most data centers a large majority of the interconnects are covered by 500 m long fibers, making pure NRZ-based transceiver as demonstrated in section 3.4.1 a more attractive solution in the search for the implementation with the lowest possible power consumption and form factor.

In both scenarios, combining 4 of these transceivers and a 4-channel Coarse Wavelength Division Multiplexing (CWDM) scheme, could be used to realize for an attractive alternative for a 400 GbE transceiver. Even though the EAM does not operate over the whole C-band, a 1-dB link power penalty bandwidth of more than 30 nm was reported for a similar, but shorter GeSi EAM in [18] . This should be more than enough to support a 4λ-CWDM configuration (i.e. for 400 GbE links), albeit located partly in (the upper) C-band and partly in (the lower) L-band. This reasoning is further validated by the successful transmissions of 100 Gb/s at 1560 nm and 1601 nm, with limited penalty.

3.4.3 100 Gb/s EAM-to-EAM Transmission

The proposed EAM is not only ideally suited as modulator, but can also function as a high-speed photodiode by increasing the reverse bias beyond the ideal modulation bias point to absorb as much light as possible. Fig.3.9 shows the eye-diagrams for different lengths of fiber in such an EAM-to-EAM link. A 40 GHz RF-probe, a 65 GHz bias-T and a 50 cm long coax-cable were used to deliver a reverse bias of 3 V to an identical copy of the EAM located on a different die. As this setup posed an additional BW-limitation, only offline BER measurements using EDB at 1560 nm were performed for which the results are depicted in Fig.3.8.

For fiber lengths up to 0.5 km the measured BERs correspond well to the BER-curves of the EAM-to-PD link. For fiber spans of 1 km, a reduction in eye height is noticeable in Fig. 3.9c, leading to slightly higher average optical power of 0.9 dBm to reach the FEC-limit (a penalty of 0.7 dB). After 2 km, the eye degradation is even more pronounced (Fig. 3.9d), but even now, sub-FEC operation is obtained above 4.1 dBm of optical input power. A similar, but smaller increase in power penalty was also observed for the

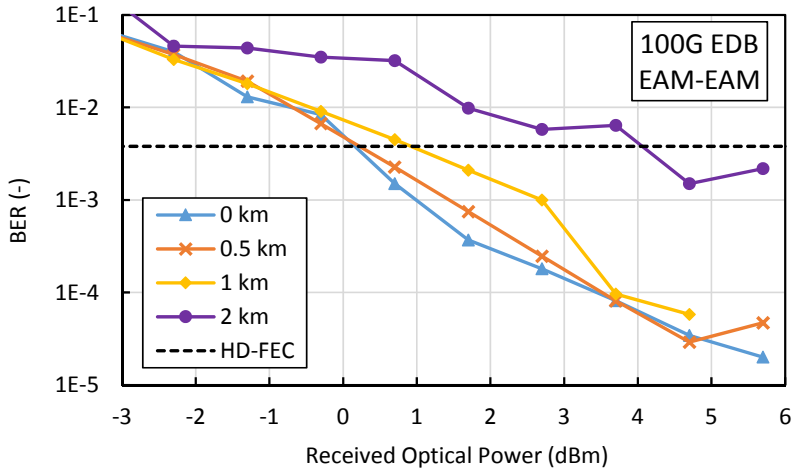

Figure 3.8: Offline measured BER curves for duobinary modulation at 1560 nm for a EAM-to-EAM link.

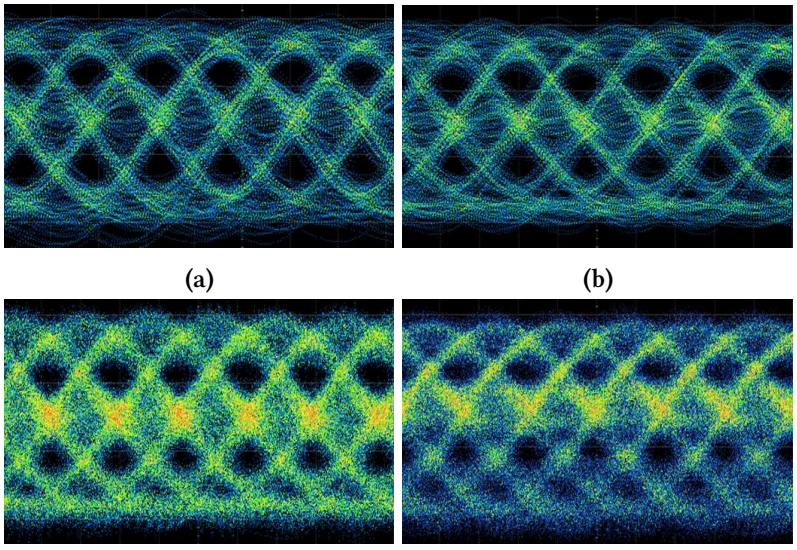

Figure 3.9: Measured eye-diagrams of 100 Gb/s EDB transmission at 1560 nm for (a) B2B, (b) 500 m, (c) 1 km, (d) 2 km of SSMF.

real-time 2 km PD-based link. This indicates that, in the presence of severe CD, additional bandwidth reductions in the E/O/E (e.g. by the bandwidth-limited input buffer of the RX-IC or by the additional 40 GHz RF-probe and 50 GHz coax-cable for the EAM-based PD) might rapidly degraded the link performance. Nevertheless, the possibility of a silicon photonics transceiver operating at line rates of 100 Gb/s based on the GeSi EAM acting as modulator and as photodetector is validated.

3.5 Conclusion

In this chapter, we have presented real-time, single-lane and serial 100 Gb/s transmission with NRZ-OOK as well as electrical duobinary using a germanium-silicon EAM in combination with an in-house developed BiCMOS-based transmitter and receiver chipset, without any need for DSP. The EAM was driven lumped without any termination with 2 V_{pp}.

For 100 Gb/s NRZ, we recorded BERs down to 6×10^{-9} in a back-to-back link, as well as successful transmissions -assuming FEC- over 500 m of SSMF and 2 km of NZ-DSF, which was comparable to 1 km of SSMF. We identified the chromatic distortion of the fiber channel as the main limitation in the link, degrading the frequency response even for relatively short fiber spans of 0.5 km to 2 km due to the high line rate. As a possible solution, a 3-level duobinary modulation scheme was investigated and verified to be much more resilient towards this effect, allowing real-time sub-FEC operation up to 2 km of SSMF.

Finally, the possibility of a silicon-based transceiver working at line rates of 100 Gb/s using the GeSi EAM both as modulator and as photodiode, was demonstrated for EDB modulation up to 2 km of SSMF. These results showcase the capabilities of silicon photonics as a possibly disruptive technology for compact and low-power transceivers for 800 GbE and 1.6 TbE short reach-optical interconnects.

References

[1] J. Verbist, M. Verplaetse, S. A. Srinivasan, J. Van Kerrebrouck, P. De Heyn, P. Absil, *et al.*, "Real-time 100 gb/s nrz and edb transmission with a gesi electroabsorption modulator for short-reach optical interconnects", *Journal of Lightwave Technology*, vol. 36, no. 1, pp. 90–96, Jan. 2018.

[2] J. Verbist, M. Verplaetse, S. A. Srivinasan, P. D. Heyn, T. D. Keulenaer, R. Pierco, *et al.*, "First Real-Time 100-Gb/s NRZ-OOK Transmission over 2 km with a Silicon Photonic Electro-Absorption Modulator", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[3] *IEEE P802.3bs 400 Gigabit Ethernet Task Force*. [Online]. Available: `http://www.ieee802.org/3/bs/`.

[4] Q. Zhang, N. Stojanovic, T. Zuo, L. Zhang, C. Prodaniuc, F. Karinou, *et al.*, "Single-lane 180 Gb/s SSB-duobinary-PAM-4 signal transmission over 13 km SSMF", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[5] T. Zuo, L. Zhang, Q. Zhang, J. Zhou, E. Zhou, and G. N. Liu, "Single lane 112-Gbps analog small form-factor pluggable module with only 4-GHz end-to-end 3-dB bandwidth employing duobinary 4-PAM", in *2016 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2016, pp. 1–3.

[6] A. Chiuchiarelli, R. Gandhi, S. M. Rossi, L. H. H. Carvalho, F. Caggioni, J. C. R. F. Oliveira, *et al.*, "Single wavelength 100G real-time transmission for high-speed data center communications", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[7] J. Lee, S. Shahramian, N. Kaneda, Y. Baeyens, J. Sinsky, L. Buhl, *et al.*, "Demonstration of 112-gbit/s optical transmission using 56gbaud pam-4 driver and clock-and-data recovery ics", in *2015 European Conference on Optical Communication (ECOC)*, Sep. 2015, pp. 1–3.

[8] T. Chan, I. C. Lu, J. Chen, and W. I. Way, "400-Gb/s Transmission Over 10-km SSMF Using Discrete Multitone and 1.3- mum EMLs", *IEEE Photonics Technology Letters*, vol. 26, no. 16, pp. 1657–1660, Aug. 2014.

[9] P. Dong, J. Lee, Y. K. Chen, L. L. Buhl, S. Chandrasekhar, J. H. Sinsky, *et al.*, "Four-Channel 100-Gb/s Per Channel Discrete Multitone Modulation Using Silicon Photonic Integrated Circuits", *Journal of Lightwave Technology*, vol. 34, no. 1, pp. 79–84, Jan. 2016.

[10] J. Lee, N. Kaneda, T. Pfau, A. Konczykowska, F. Jorge, J. Y. Dupuy, *et al.*, "Serial 103.125-gb/s transmission over 1 km ssmf for low-cost, short-reach optical interconnects", in *OFC 2014*, Mar. 2014, pp. 1–3.

[11] M. Verplaetse, R. Lin, J. V. Kerrebrouck, O. Ozolins, T. D. Keulenaer, X. Pang, *et al.*, "Real-Time 100 Gb/s Transmission Using Three-Level Electrical Duobinary Modulation for Short-Reach Optical Interconnects", *Journal of Lightwave Technology*, vol. 35, no. 7, pp. 1313–1319, Apr. 2017.

[12] H. Zwickel, T. D. Keulenaer, S. Wolf, C. Kieninger, Y. Kutuvantavida, M. Lauermann, *et al.*, "100 Gbit/s serial transmission using a silicon-organic hybrid (SOH) modulator and a duobinary driver IC", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[13] D. Patel, A. Samani, V. Veerasubramanian, S. Ghosh, and D. V. Plant, "Silicon Photonic Segmented Modulator-Based Electro-Optic DAC for 100 Gb/s PAM-4 Generation", *IEEE Photonics Technology Letters*, vol. 27, no. 23, pp. 2433–2436, Dec. 2015.

[14] A. Samani, M. Chagnon, E. El-Fiky, D. Patel, M. Jacques, V. Veerasubramanian, *et al.*, "Silicon photonics modulator architectures for multi-level signal generation and transmission", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[15] P. Groumas, Z. Zhang, V. Katopodis, and A. Konczykowska, "Tunable 100 Gbaud Transmitter Based on Hybrid Polymer-to-Polymer Integration for Flexible Optical Interconnects", *Journal of Lightwave Technology*, vol. 34, no. 2, pp. 407–418, Jan. 2016.

[16] M. Verplaetse, T. D. Keulenaer, A. Vyncke, R. Pierco, R. Vaernewyck, J. V. Kerrebrouck, *et al.*, "Adaptive Transmit-Side Equalization for Serial Electrical Interconnects at 100 Gb/s Using Duobinary", *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 64, no. 7, pp. 1865–1876, Jul. 2017.

[17] T. D. Keulenaer, G. Torfs, Y. Ban, R. Pierco, R. Vaernewyck, A. Vyncke, *et al.*, "84 gbit/s sige bicmos duobinary serial data link including serialiser/deserialiser (serdes) and 5-tap ffe", *Electronics Letters*, vol. 51, no. 4, pp. 343–345, 2015.

[18] S. A. Srinivasan, P. Verheyen, R. Loo, I. D. Wolf, M. Pantouvaki, G. Lepage, *et al.*, "50Gb/s C-band GeSi waveguide electro-absorption modulator", in *2016 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2016, pp. 1–3.

[19] S. A. Srinivasan, M. Pantouvaki, S. Gupta, H. T. Chen, P. Verheyen, G. Lepage, *et al.*, "56 Gbps Germanium Waveguide Electro-Absorption Modulator", *Journal of Lightwave Technology*, vol. 34, no. 2, pp. 419–424, Jan. 2016.

[20] J. Wang, Y. Xuan, C. Lee, B. Niu, L. Liu, G. N. Liu, *et al.*, "Low-loss and misalignment-tolerant fiber-to-chip edge coupler based on double-tip inverse tapers", in *2016 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2016, pp. 1–3.

4

DAC-less and DSP-free 112 Gb/s PAM-4 Transmitter using Two Parallel Electro-Absorption Modulators

Four-level pulse amplitude modulation (PAM-4) is the modulation format of choice for the next generation of 400 gigabit Ethernet short-reach optical transceiver. However, generating and receiving PAM-4 at line rates of 112 Gb/s has proven challenging without relying on power-hungry tools as digital signal processing and digital-to-analog converters, as it requires linearity from the E/O-components in the link and/or pre-distortion techniques. Moving the binary to multilevel conversion to the optical domain would greatly relax these requirements. Electro-absorption based transceivers would be ideally suited for this type of data center interconnects as they are capable of combining low-power and high bandwidth operation with a very compact layout, removing the need for large travelling wave structures and dedicated 50 Ω terminations.

In this chapter, we present a novel transmitter topology for generating PAM-4 using two binary driven electro-absorption modulators in parallel. Using this approach, we achieve superior performance with respect to a single, but identical multilevel driven EAM. Finally, we demonstrate the first silicon-based modulator capable of transmitting single-lane 112 Gb/s PAM-4 over

2 km of standard single-mode fiber without any electrical DAC, DSP or long transmission line structures and terminations.

This chapter is based on the journal paper with the same title that was published in Journal of Lightwave Technology (Jan. 2018) [1]. It was an invited extension of our post deadline paper presented at ECOC 2017 [2].

4.1 Introduction

Next generation transceivers for short-reach optical interconnects will likely employ a four lane scheme with 100G line rates [3], as this is a natural successor of the 100 Gigabit Ethernet (GbE) modules used today without having to increase the component and lane count and, as such, the packaging cost. Although some demonstrations of 100G line rates using non-return to zero (NRZ) or 3-level duobinary exist [4–8], four-level pulse amplitude modulation (PAM-4) has emerged as the preferred modulation format for this scenario, balancing relaxed bandwidth requirements with increased complexity for the E/O-components in the link. Currently, most of the PAM-4 transmitters at 100 Gb/s and above still require electrical digital-to-analog converters (DACs) to generate the multilevel signal [9, 10].

However, to drive a single optical modulator the DAC must provide a sufficiently large voltage swing or must be followed by a linear output driver. Both options substantially increase the power consumption of the transmitter with respect to a conventional NRZ driver at the same baud rate. Moving the DAC-operation to the optical domain would remove the linearity requirements at the transmitter, reducing the complexity of the electrical front-end and its power consumption.

Recently, several optical DACs have been proposed using a segmented Mach-Zehnder modulator (MZM) [11], parallel MZMs [12, 13], silicon ring modulators [14–16], Si-on-IIIV electro-absorption modulated distributed feedback laser (EML DFB) [17] or by using polarization division multiplexing (PDM) for the least and most significant bit (LSB and MSB) [18, 19]. Although the MZM-based solutions show good performance, they might not be suited for short-reach interconnects as they typically require large transmission line structures and dedicated, power-consuming terminations. Segmented silicon microring modulators have been used [14, 15], and recent demonstrations have shown PAM-4 transmission up to 128 Gb/s [16]. However, ring resonators are very susceptible to temperature variations and need control systems to guarantee stable

operation. In [17], direct and external modulation of a IIIV-on-Si DFB laser were combined to encode the LSB and the MSB. However, the speed was limited to 50 Gb/s PAM-4 as the bandwidth of the direct modulation (14 GHz) was a bottleneck for the overall data rate.

Another approach towards a DAC-less transmitter was shown in [18, 19], where a PDM scheme was used to transport the LSB and MSB over the optical channel with and electro-absorption modulated laser in Indium Phosphide (InP). This allows for an independent power addition at the receiver, provided it is polarization insensitive. A drawback of this method is that it already occupies both polarizations, removing the possibility of doubling the data rate by implementing a PDM scheme. Moreover, the demonstrations still rely on discrete external components to perform the polarization handling as these are not readily available in an InP integration platform. Implementing a compact, low-power 112 Gb/s PAM-4 transmitter in a Silicon platform would provide a low-cost solution, which could be produced in high volume leveraging the existing CMOS fabrication infrastructure.

In this chapter, we present a novel single-lane, single-polarization integrated PAM-4 transmitter based on the vector addition of two binary driven amplitude modulators in parallel. An integrated prototype consisting of two compact, waveguide GeSi electro-absorption modulators (EAMs) was fabricated in imec's Silicon Photonics platform, outperforming a single multilevel driven GeSi EAM. Using this prototype, we demonstrate the first transmission of 112 Gb/s PAM-4 over 2 km of standard single mode fiber with a silicon-based modulator without any DAC, DSP or large transmission line structures.

4.2 Proposed Topology

As the transfer function of an electro-absorption modulator is typically non-linear and not symmetrical (as is the case for a MZM), generating a clean PAM-4 signal with equidistant levels on a single EAM can be challenging. Often, a power-hungry DAC or some clever analog pre-distortion method is required to produce equidistant eye levels. This becomes especially difficult if the EAM has only limited extinction ratio (ER < 10 dB). Performing the multilevel generation in the optical domain rather than the electrical domain by operating the EAMs as two binary driven switches, allows us to bypass the non-linearity of the modulator. More importantly, this also eliminates the need for a DAC or linear driver

at the transmitter, allowing simple and low-power NRZ driver topologies (e.g. CMOS inverters) to be used instead. In other words, the linearity requirement is completely removed from the transmit side in both the electrical and optical domain.

Here, we present a new type of optical DAC using 2 identical, binary driven EAMs in parallel. Although we have focused on EAMs as modulators to implement the proposed optical DAC topology, any type of amplitude modulator can be used. As such, all the principles and remarks that are discussed in the following paragraphs, also apply to any choice of amplitude modulator.

4.2.1 Principle of Operation

The proposed modulator is shown in Fig. 4.2 and consists of a splitter with power ratio $\alpha : 1-\alpha$, two identical EAMs, a DC phase shifter and a 3 dB-combiner. The DC phase shift $\Delta\phi$ between both arms is needed to provide an additional degree of freedom to place the PAM-4 levels equidistantly. The input splitter can be realized as a tunable splitter (e.g. Mach-Zehnder Interferometer, using commonly available components such as 1×2 and 2×2 multi-mode interferometers) or as a custom design (e.g. a star coupler with fixed coupling ratio). When branch α corresponds to the LSB and branch $1-\alpha$ to the MSB, the output power levels P_{ij} are given by

$$P_{ij} = \frac{1}{2}\left|\frac{\sqrt{1-\alpha}}{IL \cdot ER^{1-i}} + e^{j\Delta\phi}\frac{\sqrt{\alpha}}{IL \cdot ER^{1-j}}\right|^2 \quad \text{for} \quad i,j = 0,1 \qquad (4.1)$$

Where $i, j = MSB, LSB$ and the (identical) EAMs are characterized by a bias and voltage dependent ER and insertion loss (IL). For simplicity, we assume that no phase difference is introduced between the 0 and the 1 level by the EAMs. By choosing an appropriate α, we can fix the levels where at least one EAM is absorbing (i.e. the symbols 00, 01 and 10) at an equidistant position. However, when both EAMs are transparent (i.e. generating 11) this level will always be above its equidistant position due to the cross product of both terms in (1). Adding a phase shift $\Delta\phi$ gives us an additional degree of freedom to place the 11-symbol at its equidistant position as P_{11} becomes

$$P_{11} = \frac{1}{2}IL\left(1 + 2\sqrt{(1-\alpha)\alpha}\cos(\Delta\phi)\right) \qquad (4.2)$$

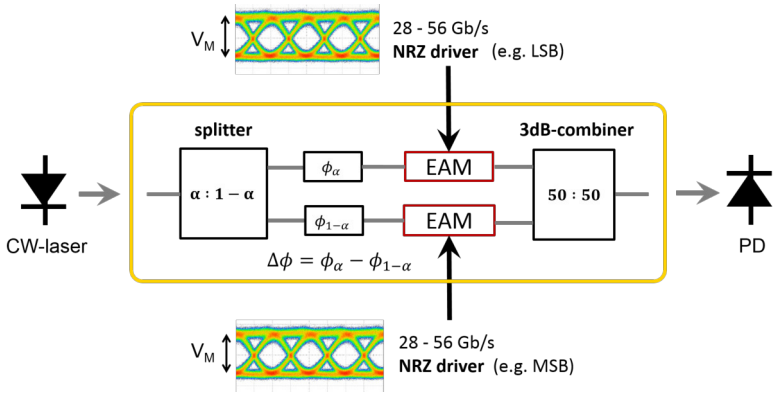

Figure 4.1: Generalized block diagram of the PAM-4 generation topology using 2 binary driven, parallel electro-absorption modulators as proposed in this work. Although the block diagram uses EAMs, this topology will work for any type of amplitude modulator.

4.2.2 Special Case: $\alpha = 1/3$ and $\Delta\phi = 90°$

A particularly interesting solution is found, when we choose $\alpha = 1/3$ and $\Delta\phi = 90°$ for which the vector diagram representing the on- and off-stated of each EAM is depicted in Fig. 4.2. In this example, the red vectors correspond to EAMs with no IL and an ER of 10 dB. This configuration has the special property that any given combination of ER, IL generates equidistant PAM-4 levels, as long as the EAMs are identical. This can be understood by realizing that both basis vectors are affected proportionally, i.e. for any choice of ER, IL the points $(\sqrt{P_{00}}, \sqrt{P_{01}}, \sqrt{P_{10}})$ always form a similar triangle, for which the ratio of its sides remains the same.

A drawback of this shaping with $\alpha = 1/3$ and $\Delta\phi = 90°$ is that the optical swing is 3 dB less than what can be maximally achieved with a single multilevel driven modulator for the same average input power, assuming full use of the available extinction by proper placement of the electrical PAM-4 levels.

Nevertheless, we will demonstrate in section IV that for modulators with a limited ER and non-linear transfer function, this penalty will be more than compensated by the improvement in eye quality.

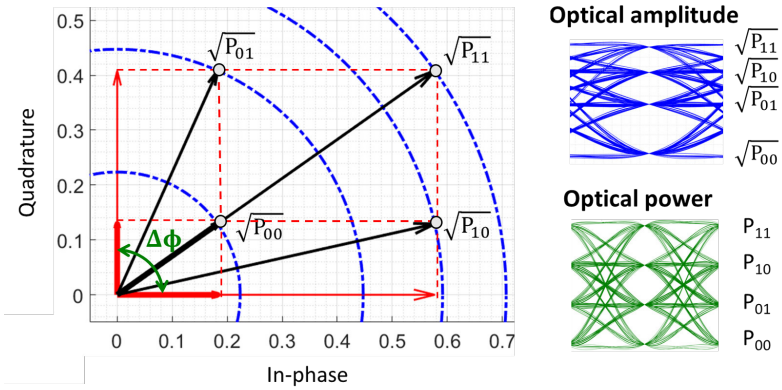

Figure 4.2: Example of an equidistant PAM-4 generation scheme, using the first quadrant of the complex plane, for the special case that the power split ratio is chosen $0.33 : 0.66$ and the phase $\Delta\phi$ between the branches is 90°. The red arrows represent the on- and off-state of the two EAMs, when driven separately. They form the basis vectors for the PAM-4 generation. The limited extinction ratio (10 dB in this example) and the resulting non-perfect zero level, is represented by the bold vectors. The four black vectors representing the four constellation points ($\sqrt{P_{00}}, \sqrt{P_{01}}, \sqrt{P_{10}}, \sqrt{P_{11}}$) are found by the vector addition of each state (on/off) of both basis vectors. Squaring the moduli of these vectors gives us the power levels of the PAM-4 signal, when received by a square-law photodiode.

4.2.3 PAM-4 Shaping by Vector Addition

Not only equidistant eyes can be obtained, but also shaped eyes (i.e. pre-distorted) can be achieved by altering the split ratio, the phase or both. Fig. 4.3 shows an examples of varying the phase (Fig3.b) or the split ratio (Fig3.c and Fig3.d) with respect to the special configuration as discussed above (Fig3.a). Equidistant eyes are not necessarily the best configuration to obtain the minimal bit error ratio (BER) using this type of transmitter, as will be discussed next. Choosing $\Delta\phi = 0°$, we lose the equidistance of the power levels but the total swing of the PAM-4 eye almost doubles, reducing this shaping IL to only 0.13 dB.

Interestingly, while the bottom and especially the top eye height increases when compared to Fig. 4.3a, the eye height of the middle eye remains identical. Thus, if the receiver is not limited by its dynamic range, the

top eye will only contribute insignificantly to the overall bit error ratio when compared to an equidistant PAM-4 eye where each eye contributes a third of the errors. This property can be exploited to improve the BER in links where this transmitter is paired with a noise limited receiver.

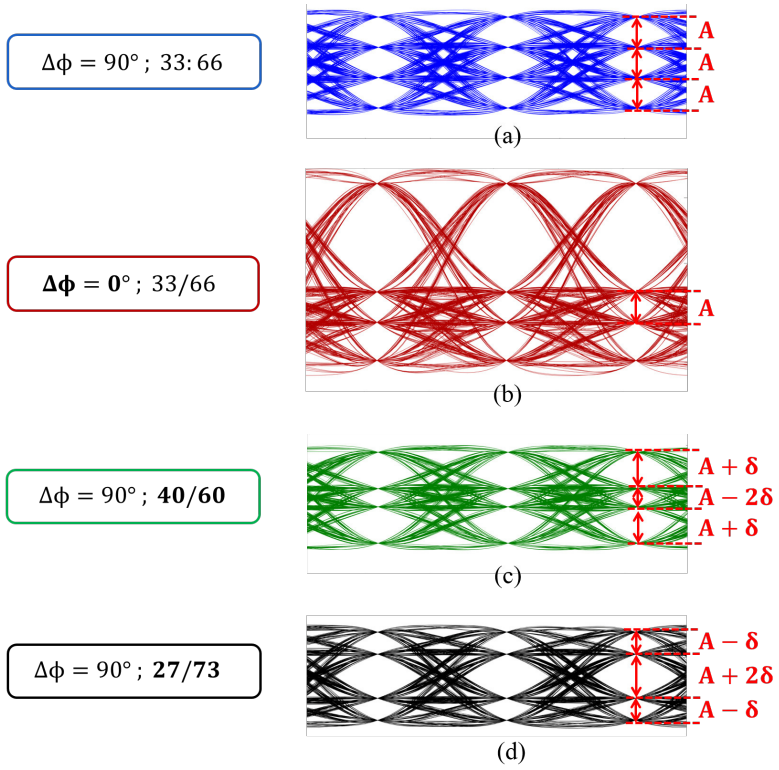

Figure 4.3: Comparison of different PAM-4 shaping though vector addition by altering the phase difference or the split ratio or with respect to for equidistant PAM-4 generation with $\Delta\phi = 90°$ and $0.33 : 0.66$ split ratio.

However, if the receiver is limited in dynamic range, we could increase α (from 33% to 40% in Fig. 4.3c) to pre-distort the multilevel signal by increasing the relative eye height of the outer eyes. This way we can compensate compression in the transimpedance amplifier (TIA) or a limited analog-to-digital converter range, relaxing the linearity requirements on the receiver frontend.

On the other hand, reducing α will introduce the inverse effect, the

inner eye height increases and the outer eye heights decrease as shown in Fig. 4.3d. This type of non-uniform PAM-4 can be beneficial in flexible passive optical networks with a spread in received optical powers to increase the aggregated capacity of the network, as was recently demonstrated in [20]. Switching between different PAM-4 shaping can be done with little extra complexity by implementing a tunable splitter (e.g. by using low-power thermal phase shifters in a Mach-Zehnder Interferometer (MZI) configuration) and a tunable phase shift between the branches (e.g. thermal phase shifter). This way, each parameter can be changed on the fly by adjusting only a single DC voltage, accommodating many different applications without having to change the transmitter design.

4.3 Experiment Setup

To verify the operation and the performance of the proposed topology, a prototype transmitter was fabricated in imec's silicon photonics platform with two standard 1×2-MMIs as splitter and combiner, a thermal heater in each arm acting as a DC phase shift and 2 identical 80 µm long GeSi EAMs. These are the same EAMs as were used in [8] and more details on a similar but shorter EAM can be found in [21, 22]. Two 50 Ω resistors are provided on-chip to allow the transmitter to be driven by an external 50 Ω-driver (RF Amp) with minimal reflection. These resistors are not necessary for the operation of the transmitter and can easily be omitted when integrated with a dedicated driver. Although the operational wavelength of the GeSi EAMs in [21, 22] red shifts approximately 0.8 nm per Kelvin due to the change in bandgap, no temperature control was needed during the experiments as these devices have a 1 dB transmitter penalty bandwidth of > 30 nm. Light is coupled in and out the photonic die through fiber-to-chip grating couplers (IL $\sim$ 6 dB/coupler).

As this structure does not have an optimized power splitting ratio between both EAMs, we mimic this effect by reducing the electrical swing on the LSB-arm with a 6 dB attenuator and by further increasing the bias voltage of the LSB-EAM. As a consequence, we suffer an additional insertion loss with respect with an optimized splitting ratio. Nevertheless, this operation allows us to validate the proposed transmitter topology.

The setup for transmission experiments is shown in Fig. 4.4 A laser source at 1577 nm with an in-fiber power of 12 dBm is coupled to the PAM-4 transmitter through fiber-to-chip grating couplers. An FPGA delivers four

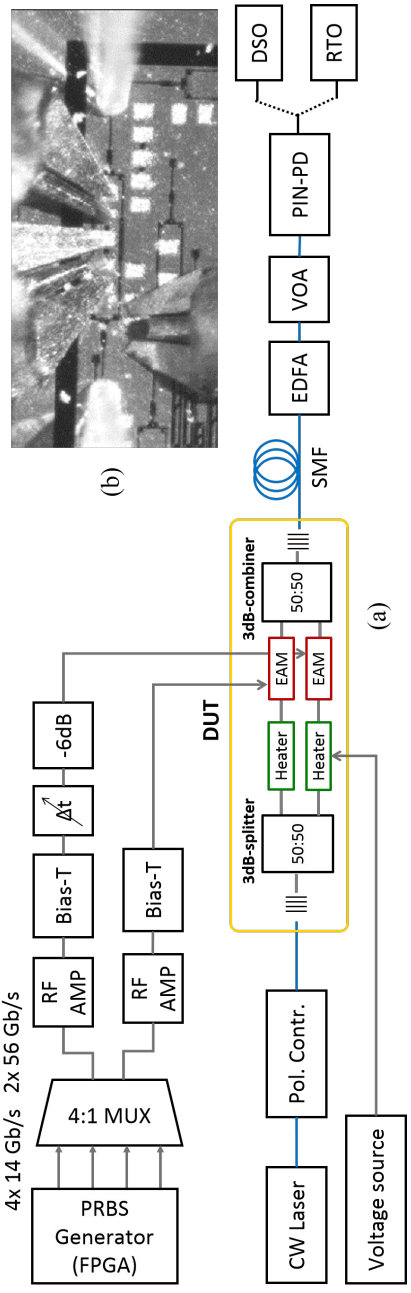

Figure 4.4: (a) Experiment setup; (b) micrograph of die during experiments.

$2^7 - 1$ long pseudo-random bit sequences (PRBS) at 14 Gb/s to an in-house developed 4-to-1 multiplexer, which generates a differential 56 Gb/s NRZ signal. To ensure decorrelation between both signals, a tunable time delay is placed after one of the differential outputs. For these first transmission experiments at 112 Gb/s only a relatively short PRBS sequence was used. Operation with longer PRBS was investigated up to 50 Gbaud with an arbitrary waveform generator (AWG) as driver. Although the performance was limited by the bandwidth of the AWG (32 GHz), no eye penalty was observed for sequence up to $2^{15} - 1$ (i.e. the longest possible PRBS that could be generated by the AWG due to its limited memory).

Next, a 50 GHz RF amplifier is added to provide a swing of $2.2\,V_{pp}$ and $1.1\,V_{pp}$ to the MSB and LSB EAM, respectively. The EAMs are biased at -0.7 V and -1.8 V through internal bias-Ts in the RF amplifier. The modulators have an estimated IL and a dynamic ER of approximately 7 dB. The average optical in-fiber power after the modulator was approximately -10 dBm. A voltage source was used to introduce a 90° phase shift between both arms. As no TIA with sufficient bandwidth (i.e. >40 GHz) was available, an erbium-doped fiber amplifier is used to compensate the insertion losses of the grating couplers and produce sufficiently large voltage swing at the output of a commercial 50 GHz photodiode (responsivity 0.65 A/W).

Although the GeSi EAMs perform slightly better around 1560 nm in terms of ER per IL [22], a longer wavelength was chosen as we only had an L-band EDFA at our disposal during the experiments. In future implementations, the EDFA can be removed from the link by incorporating a linear TIA after the photodiode and by replacing the grating couplers with low-loss edge couplers (IL < 2 dB/coupler). A variable optical attenuator (VOA) is used to fix the average input power to the photodiode to ~ 8 dBm. In the current setup, an optical modulation amplitude (OMA) of approximately 10 dBm was measured, which would correspond to an OMA close to 0 dBm in an implementation with edge-couplers (gaining $\sim$10 dB in power budget), but without the EDFA (loosing $\sim$ 20 dB in power budget). Moreover, the addition of a TIA should improve the signal-to-noise ratio in the link further by dropping the 50 Ω termination on the PD and by eliminating the amplified spontaneous noise generated in the EDFA, as there was no optical bandpass filter present in the link to minimize this noise source.

Finally, the signal is captured by a 50 GHz sampling oscilloscope (DSO) for

eye diagrams or stored by a 63 GHz 160 GSa/s real-time oscilloscope (RTO) for offline error counting. Due to the lack of a real-time PAM-4 analyzer, the BER is determined by resampling the stored signal and choosing the optimal sampling time and decision thresholds. To ensure a statistically relevant measurement, the captured waveform length was increased to produce at least 10 errors. This four-level signal is de-mapped using Gray-coding and compared to the original transmitted bit streams. No other offline DSP or equalization was used during the error counting.

4.4 Results and Discussion

The DAC operation of the prototype transmitter is verified by first driving each EAM separately in order to produce the LSB and the MSB as 56 Gb/s NRZ streams, for which the resulting optical eyes are shown in Fig. 4.5. Next, both modulators are driven simultaneously to generate the multilevel signal. The DC phase shift needed to be adjusted slightly to compensate any residual phase difference between both branches, e.g. due to an unbalanced non-zero average phase shift by operating the EAMs at different bias voltages. Nevertheless, a high-quality PAM-4 signal with clear open eyes could be generated fairly easily at 56 Gb/s and at 112 Gb/s (Fig. 4.5).

To validate the assumption that an optical DAC should have a better performance than an electrical DAC scheme as it bypasses the linearity requirements at the transmitter, a single, but identical GeSi EAM is driven with a four-level signal by a 92 GSa/s AWG. Fig. 4.6 shows the electrical input and the optical outputs for the single modulator and for the prototype transmitter (also driven by the AWG to allow a fair comparison). Even with the addition of a root-raised cosine (RRC) pulse shape by the electrical DAC, the optical DAC operation clearly outperforms a single, multilevel driven modulator.

Next, we conducted BER measurements after 0, 1 and 2 km of standard single-mode fiber at 50 and 56 GBaud, for which the received eyes and the corresponding BERs are given in Fig. 4.7. For 50 Gbaud we recorded BERs of 1.12E-6 (0 km), 4.24E-6 (1 km) and 1.4E-4 (2 km). For 56 Gbaud we obtained BERs of 1.71E-6 (0 km), 5E-5 (1 km) and 1.43E-3 (2 km). All BERs are well below the hard-decision forward error coding limit (HD-FEC) with 7% overhead of 3.8E-3, which is often used in literature to compare devices. However, in data center applications more stringent FECs apply, such as the KP-FEC (BER of 2E-4) and the KR-FEC (BER of

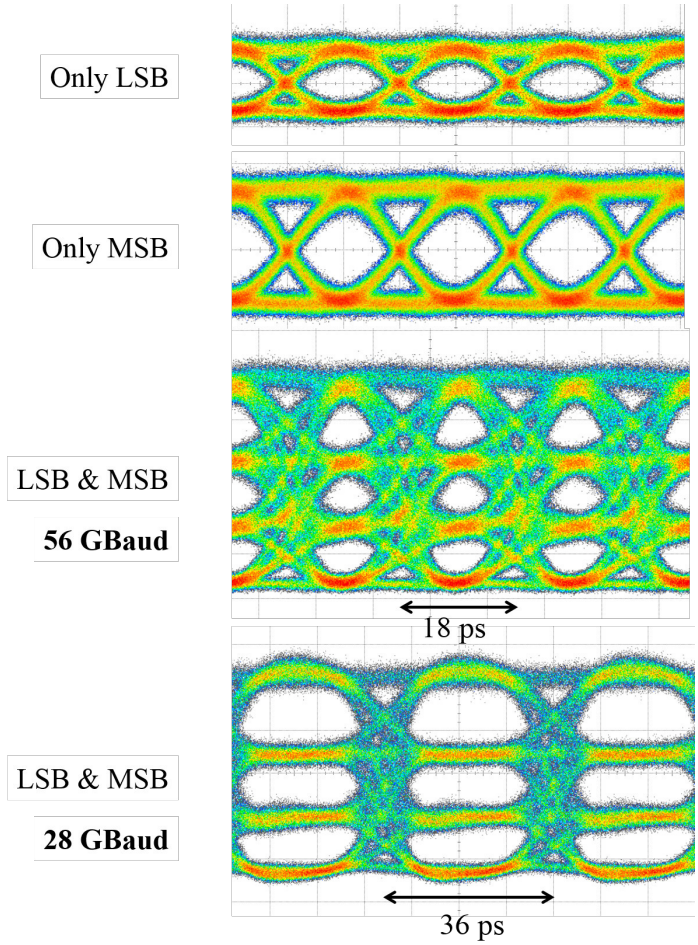

Figure 4.5: Example of the received optical eyes from the prototype transmitter with only the top or the bottom EAM driven, and with both EAMs driven at 56 Gbaud (112 Gb/s) and at 28 Gbaud (56 Gb/s) for comparison.

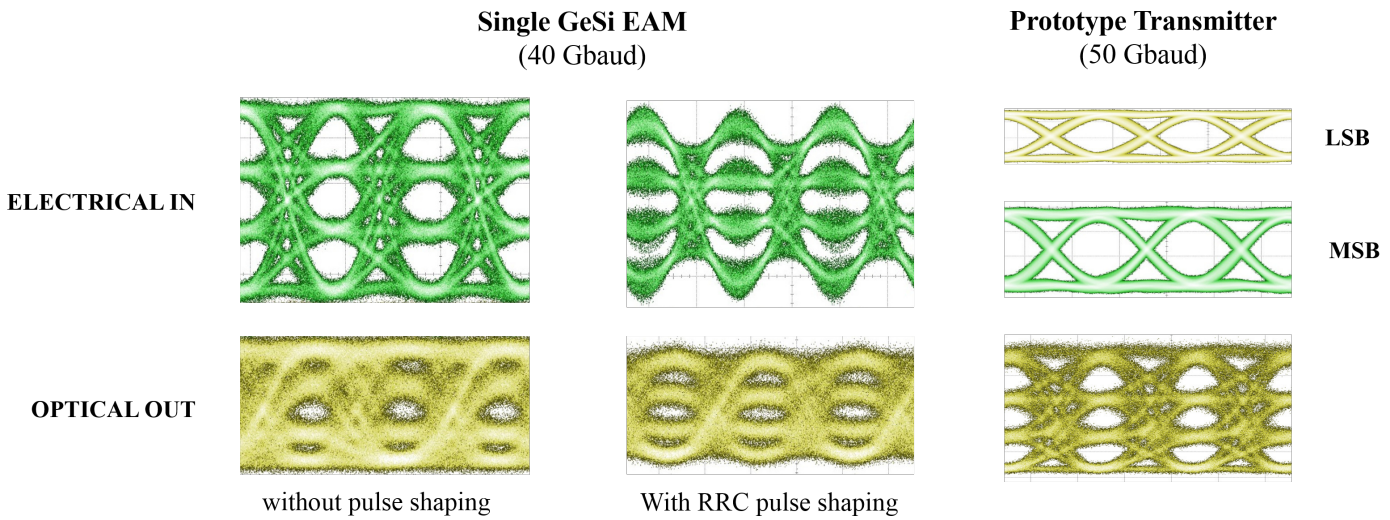

Figure 4.6: Comparison of the electrical input (top) and optical output (bottom) eyes between the multilevel driven single GeSi EAM (left) and the prototype transmitter based on the proposed topology in this work (right).

5.2E-5) [3]. Nevertheless, sub-FEC operation for both the KR- and the KP-FEC is achieved up to 1 km at 56 GBaud. At 50 GBaud, the KP-FEC can be supported up to 2 km.

The fairly large increase in BER for longer fiber spans can be largely contributed to the relatively high chromatic dispersion (CD) at 1577 nm. As we discussed in our previous work on the NRZ modulation of the EAM [8], the frequency response of the fiber channel in combination with the GeSi EAM as transmitter at 1560 nm has a frequency notch around 42 GHz for 2 km of SSMF, which is leads to a ∼2 dB penalty around 30 GHz. At 1577 nm, this notch will be at a lower frequency due to the higher CD, degrading the performance even further. Accounting for the lower CD and the improved performance of the GeSi EAMs at 1560 nm, as discussed in section III, operation below the KP-FEC limit up to 2 km should be feasible by shifting the wavelength to 1560 nm.

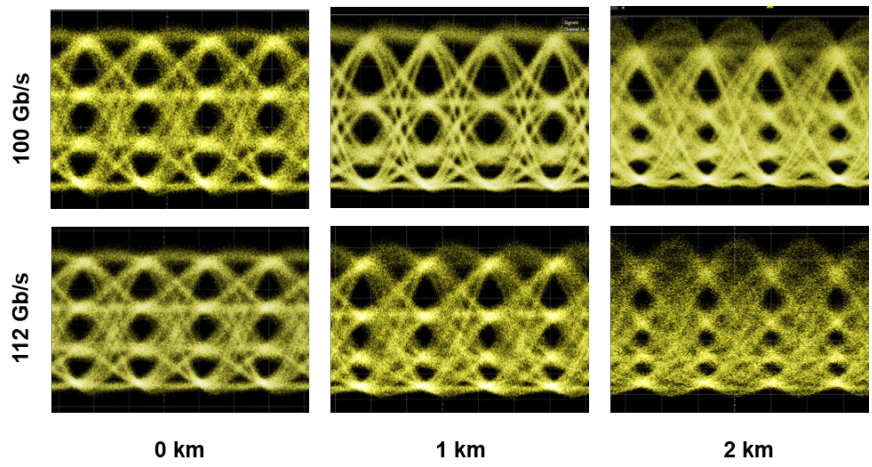

Figure 4.7: Captured eye diagrams for 50 Gbaud and 56 Gbaud PAM-4 over 0, 1 and 2 km of SSMF.

An additional benefit of using GeSi EAMs as amplitude modulators in the proposed transmitter topology, is that the same device can also be used as photodiode. Although such a link was not attempted in this experiment, we already demonstrated that these GeSi EAMs are capable of receiving 100 Gb/s NRZ with a responsivity close to 1 A/W [8].

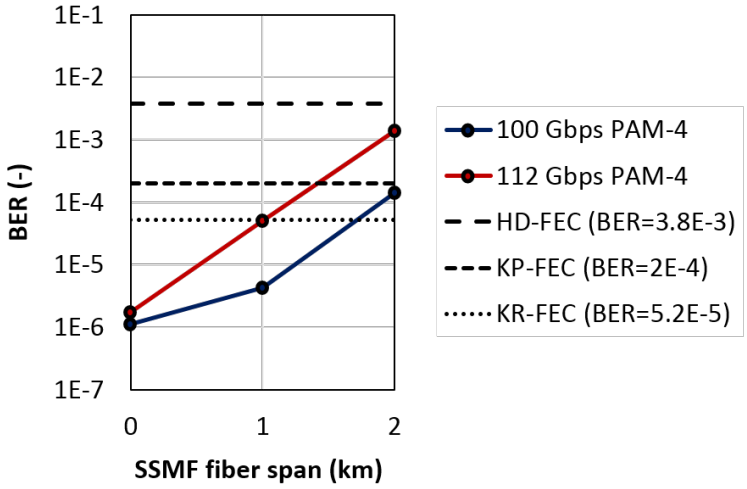

Figure 4.8: Recorded BERs for 50 Gbaud and 56 Gbaud PAM-4 over 0, 1 and 2 km of SSMF.

4.5 Conclusion

We have proposed a novel type of optical DAC to generate PAM-4, based on the vector addition of two binary driven amplitude modulators in parallel. A silicon prototype was fabricated using two GeSi EAMs, outperforming a single, multilevel driven GeSi EAM and demonstrating successful transmission over 2 km of SSMF up to 112 Gb/s. This is the first silicon-based modulator capable of generating 112 Gb/s PAM-4 without relying on power-hungry DSP, electrical DACs or long travelling wave structures and dedicated terminations. These results further showcase the benefit of postponing the DAC operation to the optical domain, as well as the bright future for SiP towards realizing compact, low-cost and low-power 400 GbE, 800 GbE and 1.6 TbE transceivers for short-reach optical interconnects.

References

[1] J. Verbist, J. Lambrecht, M. Verplaetse, J. Van Kerrebrouck, A. Srinivasan, P. De Heyn, *et al.*, "DAC-Less and DSP-Free 112 Gb/s PAM-4 Transmitter Using Two Parallel Electroabsorption Modulators", *Journal of Lightwave Technology*, vol. 36, no. 5, pp. 1281–1286, Mar. 2018.

[2] J. Verbist, J. Lambrecht, M. Verplaetse, J. Van Kerrebrouck, S. A. Srinivasan, P. De Heyn, *et al.*, "Dac-less and dsp-free pam-4 transmitter at 112 gb/s with two parallel gesi electro-absorption modulators", in *2017 European Conference on Optical Communication (ECOC)*, Sep. 2017, pp. 1–3.

[3] *IEEE P802.3bs 400 Gigabit Ethernet Task Force*. [Online]. Available: `http://www.ieee802.org/3/bs/`.

[4] J. Lee, N. Kaneda, T. Pfau, A. Konczykowska, F. Jorge, J. Y. Dupuy, *et al.*, "Serial 103.125-gb/s transmission over 1 km ssmf for low-cost, short-reach optical interconnects", in *OFC 2014*, Mar. 2014, pp. 1–3.

[5] M. Verplaetse, R. Lin, J. V. Kerrebrouck, O. Ozolins, T. D. Keulenaer, X. Pang, *et al.*, "Real-Time 100 Gb/s Transmission Using Three-Level Electrical Duobinary Modulation for Short-Reach Optical Interconnects", *Journal of Lightwave Technology*, vol. 35, no. 7, pp. 1313–1319, Apr. 2017.

[6] H. Zwickel, T. D. Keulenaer, S. Wolf, C. Kieninger, Y. Kutuvantavida, M. Lauermann, *et al.*, "100 Gbit/s serial transmission using a silicon-organic hybrid (SOH) modulator and a duobinary driver IC", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[7] P. Groumas, Z. Zhang, V. Katopodis, and A. Konczykowska, "Tunable 100 Gbaud Transmitter Based on Hybrid Polymer-to-Polymer Integration for Flexible Optical Interconnects", *Journal of Lightwave Technology*, vol. 34, no. 2, pp. 407–418, Jan. 2016.

[8] J. Verbist, M. Verplaetse, S. A. Srinivasan, J. Van Kerrebrouck, P. De Heyn, P. Absil, *et al.*, "Real-time 100 gb/s nrz and edb transmission with a gesi electroabsorption modulator for short-reach optical interconnects", *Journal of Lightwave Technology*, vol. 36, no. 1, pp. 90–96, Jan. 2018.

[9] A. Chiuchiarelli, R. Gandhi, S. M. Rossi, L. H. H. Carvalho, F. Caggioni, J. C. R. F. Oliveira, *et al.*, "Single wavelength 100g real-time transmission for high-speed data center communications", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[10] Q. Zhang, N. Stojanovic, T. Zuo, L. Zhang, C. Prodaniuc, F. Karinou, *et al.*, "Single-lane 180 gb/s ssb-duobinary-pam-4 signal transmission over 13 km ssmf", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[11] D. Patel, A. Samani, V. Veerasubramanian, S. Ghosh, and D. V. Plant, "Silicon Photonic Segmented Modulator-Based Electro-Optic DAC for 100 Gb/s PAM-4 Generation", *IEEE Photonics Technology Letters*, vol. 27, no. 23, pp. 2433–2436, Dec. 2015.

[12] A. Samani, V. Veerasubramanian, E. El-Fiky, D. Patel, and D. V. Plant, "A silicon photonic pam-4 modulator based on dual-parallel mach–zehnder interferometers", *IEEE Photonics Journal*, vol. 8, no. 1, pp. 1–10, Feb. 2016.

[13] M. Chagnon, T. A. Eriksson, R. Dischler, F. Buchali, H. Bülow, and S. ten Brink, "Duobinary iq modulation schemes for c- and o-band pam4 direct detect systems", in *2017 European Conference on Optical Communication (ECOC)*, Sep. 2017, pp. 1–3.

[14] R. Li, D. Patel, A. Samani, E. El-Fiky, Z. Xing, M. Sowailem, *et al.*, "An 80 gb/s silicon photonic modulator based on the principle of overlapped resonances", *IEEE Photonics Journal*, vol. 9, no. 3, pp. 1–11, Jun. 2017.

[15] R. Li, D. Patel, A. Samani, E. El-Fiky, Z. Xing, M. Morsy-Osman, *et al.*, "Silicon photonic ring-assisted mzi for 50 gb/s dac-less and dsp-free pam-4 transmission", *IEEE Photonics Technology Letters*, vol. 29, no. 12, pp. 1046–1049, Jun. 2017.

[16] J. Sun, R. Kumar, M. Sakib, J. B. Driscoll, H. Jayatilleka, and H. Rong, "A 128 Gb/s PAM4 Silicon Microring Modulator With Integrated Thermo-Optic Resonance Tuning", *Journal of Lightwave Technology*, vol. 37, no. 1, pp. 110–115, Jan. 2019.

[17] A. Abbasi, L. Abdollahi Shiramin, B. Moeneclaey, J. Verbist, X. Yin, J. Bauwelinck, *et al.*, "Iii–v-on-silicon c-band high-speed electro-absorption-modulated dfb laser", *Journal of Lightwave Technology*, vol. 36, no. 2, pp. 252–257, Jan. 2018.

[18] W. Huang, C. Wei, and J. Chen, "Optical dac for generation of pam4 using parallel electro-absorption modulators", in *ECOC 2016; 42nd European Conference on Optical Communication*, Sep. 2016, pp. 1–3.

[19] M. Theurer, H. Zhang, Y. Wang, W. Chen, L. Chen, Zeng, *et al.*, "2 × 56 gb/s from a double side electroabsorption modulated dfb laser and application in novel optical pam4 generation", *Journal of Lightwave Technology*, vol. 35, no. 4, pp. 706–710, Feb. 2017.

[20] R. van der Linden, N. C. Tran, E. Tangdiongga, and A. M. J. Koonen, "Improvement on received optical power based flexible modulation in a pon by the use of non-uniform pam", in *2017 European Conference on Optical Communication (ECOC)*, Sep. 2017, pp. 1–3.

[21] M. Pantouvaki, S. A. Srinivasan, Y. Ban, P. De Heyn, P. Verheyen, G. Lepage, *et al.*, "Active components for 50 gb/s nrz-ook optical interconnects in a silicon photonics platform", *Journal of Lightwave Technology*, vol. 35, no. 4, pp. 631–638, Feb. 2017.

[22] S. A. Srinivasan, P. Verheyen, R. Loo, I. De Wolf, M. Pantouvaki, G. Lepage, *et al.*, "50gb/s c-band gesi waveguide electro-absorption modulator", in *2016 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2016, pp. 1–3.

5

Real-Time and DSP-free 128 Gb/s PAM-4 Link using a Binary Driven Silicon Photonic Transmitter

In this chapter, we continue with the optical PAM-4 tranmitter topology based on the interferometric combination of two GeSi EAMs as introduced in the previous chapter. We increase the baudrate to 64 Gbaud and try to establish a link in real-time. As no real-time bit-error rate testers (BERT) above 53 Gbaud were not available at the time of the experiments, we built are own custom real-time PAM-4 BERT by combining two electrical duobinary receivers that were already used in chapter 3.

Using this setup, we demonstrate the first real-time 128 Gb/s PAM-4 transmission with a silicon photonic transmitter in a chip-to-chip link. In a back-to-back scenario, we obtained a bit-error ratio (BER) of 4×10^{-10} *without requiring any DAC, DSP, or modulators with large traveling wave structures. Over 1 km of standard single mode fiber a BER of* 8×10^{-6} *is recorded, still well below the KP4 forward error-coding limit. These results correspond to the lowest BERs reported for any real-time PAM-4 link at 100 Gb/s or higher at the time of publishing, illustrating the benefit of performing the DAC operation in the optical domain. Additionally, we elaborate on the advantages and disadvantage of inducing the power difference between the LSB and the MSB*

electrically or optically.

This chapter is based on the invited journal paper with the same title published in the Journal of Lightwave Technology (Jan. 2019) [1].

5.1 Introduction

The adoption of the 400 Gigabit Ethernet (GbE) standards has made four-level pulse amplitude modulation (PAM-4) the modulation format of choice for the next-generation single-mode data center interconnects (DCI). The 400GBASE standard specifies 8 lanes of 53.125 Gb/s PAM-4 for 1 km and 2 km standard single-mode fiber (SSMF) links and 4 lanes of 106.25 Gb/s PAM-4 for 500m SSMF links, introducing the first 100 Gb/s per wavelength standard [2].

However, also for longer fiber spans a $4 \times 100G$ PAM-4 scheme is a likely candidate, as it offers the lowest practical lane count and thus the most compact transceiver. Regardless of the outcome of these ongoing standardizations, the wide spread deployment of 100 Gb/s/λ modules remains a logical step on the growth path of data centers, irrespective of the interconnect span. Possibly an even more challenging task is updating the copper interconnects to sustain these rapidly increasing data rates. Moreover, with the increasing data rates next-generation optical interconnects are expected to move to the intra-rack and intra-board interconnects [3]. Especially for these on-board optical interconnects, minimizing power and area (both electrically and optically) will be of the utmost importance.

Previously, several examples of 100 Gb/s/λ PAM-4 links have been demonstrated [4–14]. However, many of these examples had to rely on DACs, ADCs, and/or digital signal processing (DSP) at the transmitter and/or at the receiver [4–9], leading to a significant increase in latency, power consumption, and cost. Often the required DSP at the receiver prevents online (or real-time) link experiments, even with high-end test equipment.

Nevertheless, some real-time examples demonstrating 100 Gb/s single-lane PAM-4 transmission have been reported [10, 11]. In [10], real-time 56 Gbaud PAM-4 transmission on a discrete LiNbO3 Mach-Zehnder modulator (MZM) was reported. The first real-time demonstration using a polymer on silicon MZM at 53.125 Gbaud was shown in [11], with online DSP. Both experiments employ large travelling wave modulators

(at least several millimeters long) and consequently need to be electrically terminated (typically with a 50 Ω resistor), consuming a significant amount of power and transceiver real-estate.

Recently, several DAC-less modulators have been reported, typically using either serial (or segmented) or parallel modulators. These modulators only require binary electrical signals to generate an optical multilevel signal, removing the need for power-hungry DACs and/or linear drivers. Above 50 Gbaud the integrated examples are limited to travelling-wave MZMs [8] or segmented microring modulators [9]. As previously discussed, MZMs might be ill-suited for data center applications due to their size and power consumption. Microring modulators on the other hand are very compact, but also very susceptible to temperature variations, requiring additional control systems to guarantee stable operation.

In this chapter, we use a novel optical PAM-4 generator based on two binary driven GeSi electroabsorption modulators (EAMs) in an interferometer topology, as was introduced in Chapter 4. As the EAMs are only 120 µm long, they can be driven lumped without any travelling wave electrodes or terminations. Combining this modulator with an in-house developed transceiver chipset, we are able to demonstrate the first real-time, single-wavelength transmission of 128 Gb/s PAM-4 in a chip-to-chip link. Bit-error ratios (BERs) comfortably below the KP4 forward error coding (FEC) limit of 2.4×10^{-4}, as is commonly used in data center interconnects [2], are obtained for spans up to 1 km of SSMF, without requiring any power-hungry DACs, ADCs or DSP.

5.2 Optical PAM-4 Generation

As all high-speed optical modulators are characterized by a non-linear transfer function, the most straightforward solution has been to compensate these non-idealities in the electrical domain by predistorting the levels of the applied PAM-4 signal and/or by equalizing the frequency response of the electro-optical channel. Although this solution can be very effective in leveraging non-ideal electrical and optical components, it comes at a substantial increase in the total power consumption, size, and complexity of the transceiver.

Rather than allocating transceiver resources in linearizing the electrical-to-optical conversion of a single intensity modulator, we propose to postpone the DAC operation until the optical domain by using two parallel

intensity modulators [12]. This removes the linearity requirements at the transmit side, both in the optics and the electronics. Adding a second modulator means we also need to provide a second modulator driver. Still, two NRZ drivers are likely to be more power efficient than one multilevel driver. Both drivers can be designed for non-linear operation, allowing other driver topologies to be considered (e.g. inverters) and maintaining compatibility with CMOS-based electronics and all the advantages that come with it.

In Chapter 4, we proposed a topology consisting of two intensity modulators in interferometer with 90° phase difference, where the LSB/MSB coding is realized through a 33:66 power ratio between both branches (Fig. 5.1a) [12]. Fig. 5.1b provides a more detailed explanation of the optical DAC operation using a vector diagram drawn in the first quadrant of the complex plane. The EAMs are 120µm long GeSi-based devices fabricated on imec's 200 mm platform. Their operation is based on the Franz-Keldysh effect. The bandgap of the material shifts when an electrical field is applied, changing its optical absorption spectrum. More information about on these GeSi devices can be found in [13, 14]. EAMs have the advantage that they can be made very short (minimizing the modulator capacitance and allowing them to be driven as a small lumped capacitor), removing the need for long transmission lines and power-consuming resistive terminations. However, to be able to drive the EAMs with commercially available 50 Ω-drivers (i.e. RF amplifiers), we have placed 50Ω resistors between the bondpads of each modulator to provide a matched interface on the PIC. These resistors are not necessary for the operation of the transmitter and can easily be omitted when integrated with a dedicated driver. Moreover, we recently demonstrated that such a dedicated driver for these type of EAMs can be made extremely power efficient, consuming only 61 mW at 70 Gb/s or less than 0.9 pJ/bit [15].

These first generation devices are implemented with standard 1x2 multimode interferometers (MMI) with equal power split, both at the input and at the output of the interferometer. Therefore, the LSB/MSB weighing between both branches can be achieved by reducing the voltage swing to the LSB-EAM, as illustrated in Fig. 5.2. In the experiment setup, shown in Fig. 5.3, a 6 dB attenuator was placed after one of the RF amplifiers. Due to the non-linear characteristics of the EAMs this does not perfectly correspond to a 3 dB lower optical modulation amplitude in the LSB arm. Additional corrections are applied by slightly adjusting the bias voltage of the LSB-EAM. However, this emulation only approximates the

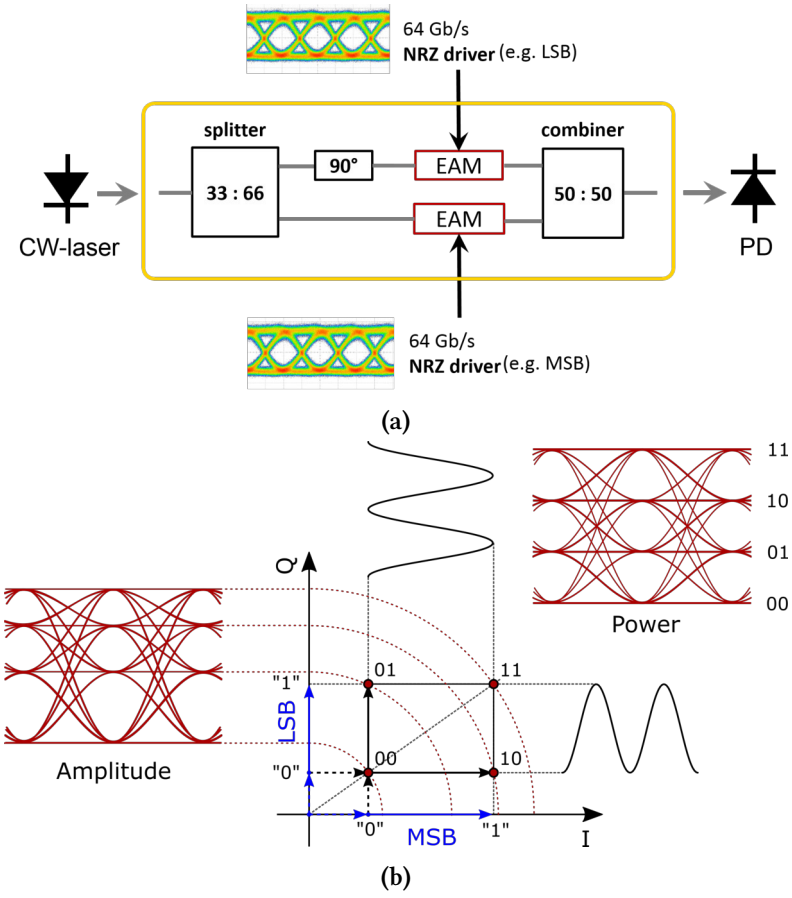

Figure 5.1: (a) 2-bit optical DAC consisting of two intensity modulators (EAMs). (b) Vector and eye diagrams of the proposed topology optical PAM-4 generator. The blue vectors represent the on- and off-state of the two EAMs, when driven separately (assuming for simplicity that no phase difference is introduced between the 0 and the 1 level by the EAMs). They form the basis vectors for the PAM-4 generation and realize a rectangular constellation in the upper quadrant of the complex plane. The limited extinction ratio (ER of 10 dB in this example) and the resulting non-perfect zero level, are represented by the dotted vectors. As long as the 90° angle is preserved between both vectors, the ER does not influence the relative positioning of the power levels, only the maximal modulation depth of the PAM-4. Even if the EAMs behave as non-perfect switches (limited ER, unbalanced IL, non-zero average phase-shift), PAM-4 can still properly be generated by adjusting the phase and/or power split. Vice versa, these parameters can also be used to predistort the PAM-4 levels (e.g. to compensate for compression in the receiver). A more in-depth discussion can be found in [12].

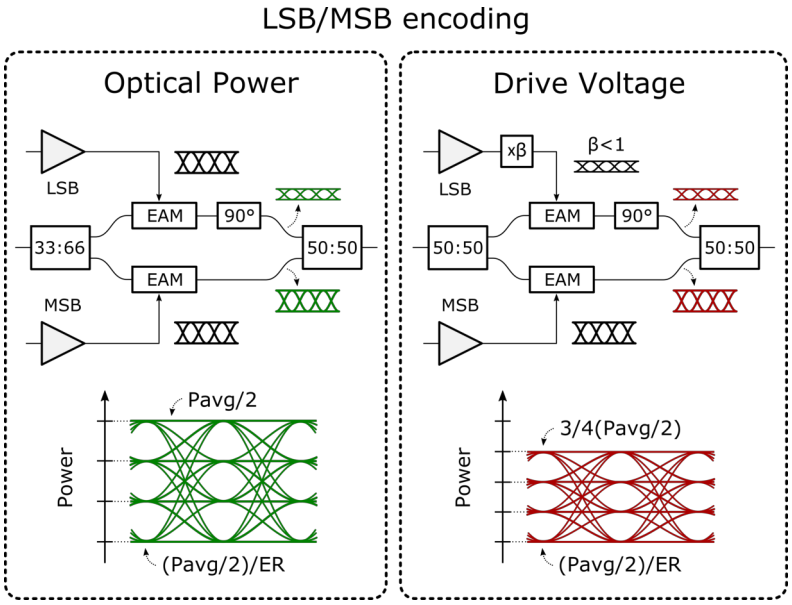

Figure 5.2: Two different versions of the optical PAM-4 generator using two parallel EAMs in an interferometer with90° phase difference: (left) using an unequal optical power split and (right) using an unequal drive voltage. To achieve equidistant PAM-4 with either of these topologies, the OMA (optical modulation amplitude) of the MSB needs to be twice the OMA of the LSB, when driven separately. This condition translates to (left) a $66:33$ power ratio or (right) a $1:\beta$ voltage swing ratio ($\beta < 1$, with $\beta = 0.5$ for a linear EAM) between the MSB and the LSB EAM. Although both version generate equidistant PAM-4, they do not have the same efficiency. Comparing the eye diagrams produced by both versions, we see that the optical power split always produces a OMA that is 33% better (or 1.25 dB higher) than the OMA of the voltage weighted version, for a given modulator and average input power (Pavg).

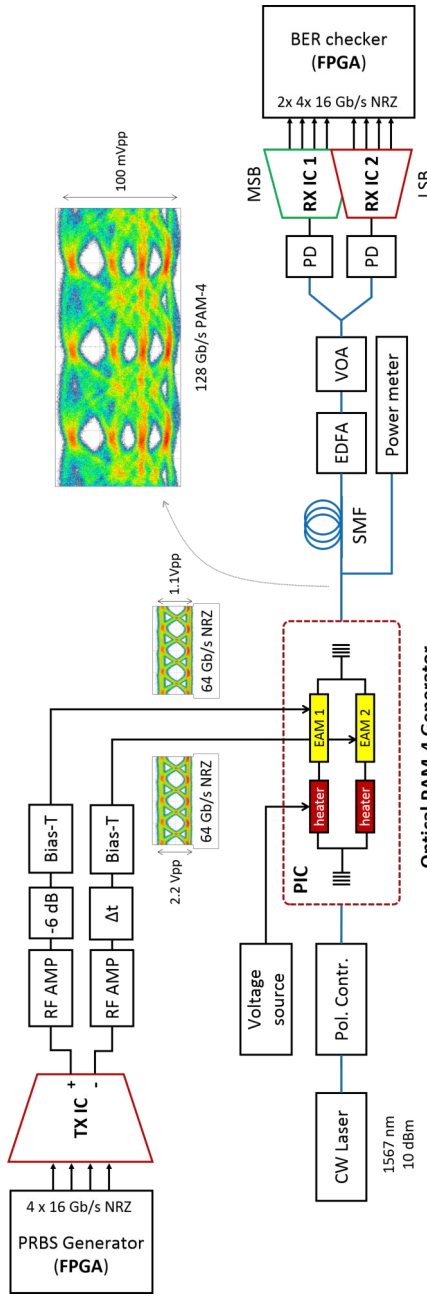

Figure 5.3: Experiment setup for the 128 Gb/s PAM-4 link.

intended PAM-4 generation, resulting in slightly unequal PAM-4 levels and reducing the overall performance. With a 33:66 power split the transmitter eye levels would be perfectly equidistant if both modulators were driven with the same voltage swing, no matter their specific transfer function (assuming no parasitic phase shift is introduced between the 0 and the 1 by the EAM) [12]. This 33:66 power split can easily be accomplished by replacing one of the MMIs with a tunable Mach-Zehnder interferometer (e.g. consisting of two MMIs and a thermal phase shifter). Moreover, Fig. 5.2 also shows that the version with unequal power split will always be more efficient than version with unequal drive voltage version. For a given modulator, maximum drive voltage and optical input power, the power split version will produce PAM-4 with a 1.25 dB higher OMA. A micrograph of the used photonic IC (PIC) with equal power split can be seen in Fig. 5.4.

An in-house developed transmitter (TX-IC in Fig. 5.4) is used to multiplex four 16 Gb/s pseudo-random bit streams (PRSB), originating from the FPGA, into one serial 64 Gb/s $2^9 - 1$ long PRBS signal. After the MUX, an analog 6-tap feed-forward equalizer can be set to improve the frequency response of the following components in the channel. In our experiments, it was mainly used to compensate the strong frequency roll-off of the electrical receiver and the RF amplifier. The differential outputs of the TX-IC are decorrelated with a mechanically tunable time delay to provide independent 64 Gb/s NRZs streams to each modulator. Finally, the signals are externally amplified to 1.1 V_{pp} and 2.2 V_{pp} for the LSB and the MSB EAM, respectively. The voltage swing could be slightly higher than in the previous experiments on similar devices [12] as the device is permitted to operate non-linearly, allowing the EAMs to be operated as switches for maximal ER. The EAMs were biased at -0.8V and -0.55V for the MSB and the LSB arm, respectively. The heater bias voltage was approximately 7V, drawing 1.5 mA. However, a fixed 90° phase difference without any power consumption can be easily realized, if we change the input splitter from a 1x2 MMI to a 2x2 MMI.

Light from a 10 dBm laser at 1567nm is coupled in and out of the PIC though fiber-to-chip grating couplers with an insertion loss of approximately 5.5 dB/coupler. The EAMs have an estimated insertion loss of 8 dB and a dynamic extinction ratio of approximately 10 dB, resulting in an average in-fiber power around -10 dBm during operation. No temperature control was used during the experiments.

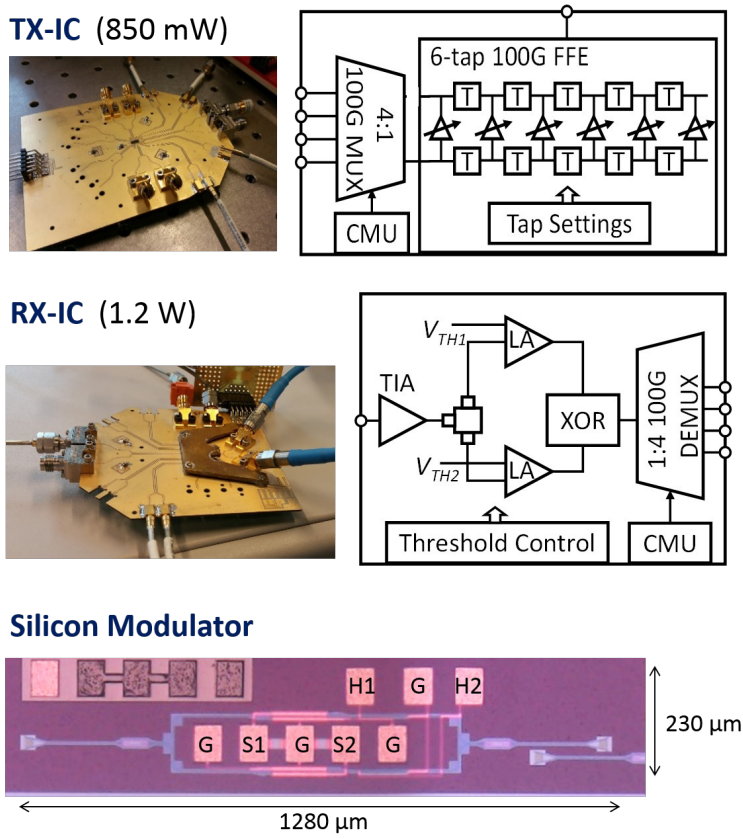

Figure 5.4: Photograph and block diagram of the BiCMOS transmitter IC (TX-IC), the BiCMOS receiver IC (RX-IC), and an annotated die micropgraph of the silicon photonic modulator. More details about the TX-IC and RX-IC can be found in [16].

5.3 Real-Time PAM-4 BER Tester

Using a transmitter as discussed in the previous section allows us to generate open PAM-4 eyes up to 64 Gbaud without any DSP, thanks to the optical DAC topology. However, receiving these signals in real-time is particularly challenging as a commercial 64 Gbaud PAM-4 bit-error rate tester (BERT) did not yet exist at the time of the experiments. Real-time oscilloscopes are not an option as the captured data has to be saved and processed offline. Still, even with a fast enough BERT it is often difficult to estimate what the performance of the link would be without any high-end

test equipment.

Therefore, we implemented a custom electrical receiver using an in-house developed chip (RX-IC) that was originally designed to decode a 3-level duobinary signal [16]. Fig. 5.4 shows the block diagram of the RX-IC as well as a photograph of the die mounted on a high-speed printed circuit board. The receiver consists of two comparators to monitor the upper and lower triangular duobinary eyes. Their outputs are XOR-ed to reproduce the original binary data and finally deserialized to a quarter-rate NRZ signal.

Fig. 5.5 explains how two of these receivers can be used to obtain both the most and least significant bit (MSB and LSB) of the PAM-4 signal simultaneously. This is in contrast to the few currently available solutions, where each of the three eyes is evaluated sequentially and the three BERs are averaged. The received PAM-4 signal is split 50:50 and fed to two photodetectors, one for each RX-IC. Both photodetectors have a similar responsivity (0.6 A/W) and a flat frequency response (important to prevent any signal distortion when converting to the electrical domain) up to 40 and 50 GHz.

Alternatively, we could have opted to split the signal after the conversion in the electrical domain, saving one photodiode. However, as the RX-IC only provides a demultiplexed and retimed output, the second photodiode could be used to monitor the eye diagram during the optimization of the equalizer settings for optimal BER (i.e. the BER of the MSB or the LSB depending on which photodiode is being monitored). As no linear high-speed transimpedance amplifier (TIA) was available, an erbium-doped fiber amplifier (EDFA) was used to provide a sufficiently large signal at the inputs of the RX-ICs. The EDFA can be omitted by replacing the relatively inefficient grating couplers with low-loss edge couplers (< 2 dB/coupler) and adding a TIA.

Decoding the MSB

Decoding the MSB is relatively straightforward, as the duobinary receiver can be easily reduced to a conventional binary receiver by setting the threshold V_{TH1} to its highest possible level, making the XOR operation transparent for the second comparator. The MSB is then received by centering V_{TH2} around the DC-level, looking into the middle eye.

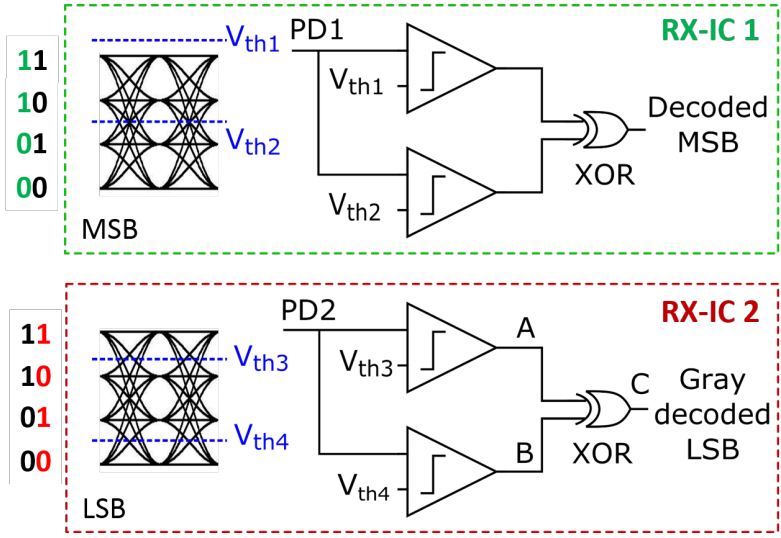

PAM-4 symbol	Gray coded **inputs**	Received MSB	A	B	C LSB	Gray decoded **outputs**
00	00	0	0	0	0	00
01	01	0	0	1	1	01
10	11	1	1	1	0	10
11	10	1	0	1	1	11

Figure 5.5: Implementation of the real-time PAM-4 receiver with automatic Gray code demapper and bit-error rate tester (BERT).

Decoding the LSB

To decode the LSB, each comparator of RX-IC 2 is set to the middle of one of the outer eyes. However, the outer eyes by themselves do not provide sufficient information to decide the received LSB. It is instructive to think of a duobinary receiver as a device determining if a symbol transition happened inside (delivering a 1) or outside (delivering a 0) both comparator thresholds. Combining this information with that of the MSB (i.e. has a symbol transition happened in the upper or lower half of the eye), resolves exactly which eye of the three has seen a symbol transition and therefore the LSB. Lastly, to obtain the LSB we need to

XOR it once more with the MSB. However, we purposefully omit this operation as the receiver now also performs an automatic demapping of a Gray encoded PAM-4 symbol, making this an additional advantage of this receiver topology, as Gray mapping and demapping is a required action set by most data center transceiver standards. Furthermore, in a test setup one can choose to Gray code the transmitted data or not, as this does not affect the BER performance when transmitting PRBS data. As both the MSB and the LSB are essentially two equal length PRSB streams, XOR-ing both will produce again an equivalent PRBS stream on which the BER counter (implemented on the FPGA) can lock.

5.4 Results and Discussion

The stand-alone performance of the transmitter is first assessed by connecting one of the photodiodes directly to a 50 GHz sampling oscilloscope and tuning the equalizers settings towards best eye quality. The heater was detuned slightly to produce a smaller phase difference, increasing the OMA of the upper eye without reducing that of other eyes as described in [12]. As long as the receiver is not limited in dynamic range, this effect can be used to improve the BER.

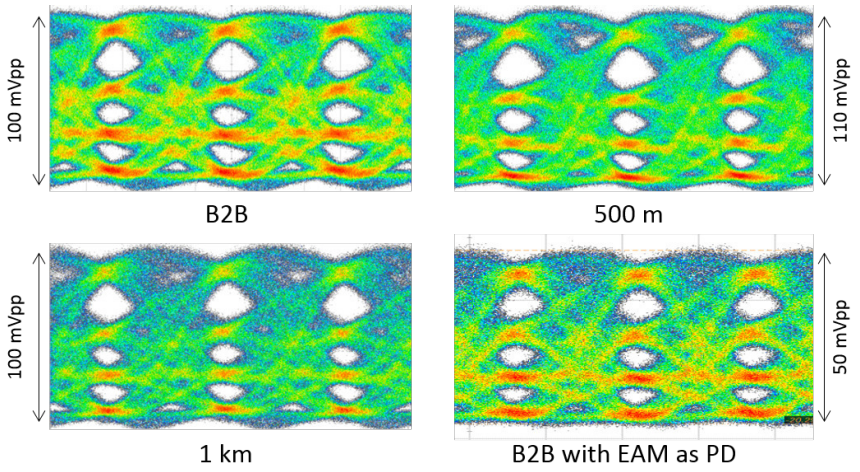

Figure 5.6: Eye diagrams at 64 Gbaud in the back-to-back case and after transmission over 500 m and 1 km of standard single mode fiber. An eye diagram with a GeSi EAM used as photodetector (back-to-back) is also shown.

Fig. 5.6 shows the eye diagrams in the back-to-back (B2B) case and after transmission over 500 m and 1 km of standard single mode fiber. Clear open eyes are obtained for all links. Next, the photodiode is connected again to the RX-IC to evaluate the real-time BER performance. The FFE parameters have to be re-optimized to minimize the BER, indicating that frequency response of the RX-ICs indeed limits the link performance. Furthermore, as the receiver was designed for 3-level duobinary reception, it did not need to be extremely linear. For example a symmetrical, gradually saturating transfer function would suffice to decode duobinary without experiencing large sensitivity penalties.

Nevertheless, for PAM-4 almost error-free operation was obtained in a back-to-back link with BERs down to 4×10^{-10}. For transmission over 500m and 1 km of SSMF we measured BERs of 7×10^{-9} and 8×10^{-6}, respectively, which is well below the KP4-FEC limit of 2.4×10^{-4} commonly used in data center interconnects, as shown in Fig. 5.7. These BERs are to the lowest reported values for a real-time PAM-4 link above 50 Gbaud [10, 11]. This affirms once more the performance benefit of the DAC-less solution over a single multilevel driven modulator as was previously observed for the modulator used in this experiment [12], as well as for other optical DACs such as segmented MZMs [8] and polarization multiplexed EAMs [17].

5.4.1 Using the GeSi EAM as a photodiode

Apart from their compact form factor and high bandwidth, the EAMs can also be used as photodetectors by setting them to maximal absorption. In [16, 18], we already observed that the EAMs can be used as optical receivers for 100 Gb/s NRZ and 3-level duobinary with a high responsivity (~ 0.8A/W for 80 µm devices). However, to receive PAM-4, the frequency response should ideally be as flat as possible. We can compare the performance by switching the transmitted eye between the photodiode with a known flat response beyond 50 GHz (upper left eye diagram in Fig. 5.6) and one of the transmitter EAMs biased at -3V on a separate die (lower right eye diagram in Fig. 5.6). Apart from some additional noise, likely due the added insertion loss of the extra grating coupler and 3 dB-splitter, little or no additional signal degradation is observed. This validates that the EAMs not only have a high responsivity but also a flat frequency response up to at least 50 GHz, making them well-suited for receiving multilevel signals. Having a single active high-speed optical component for both the transmitter and the receiver, would greatly

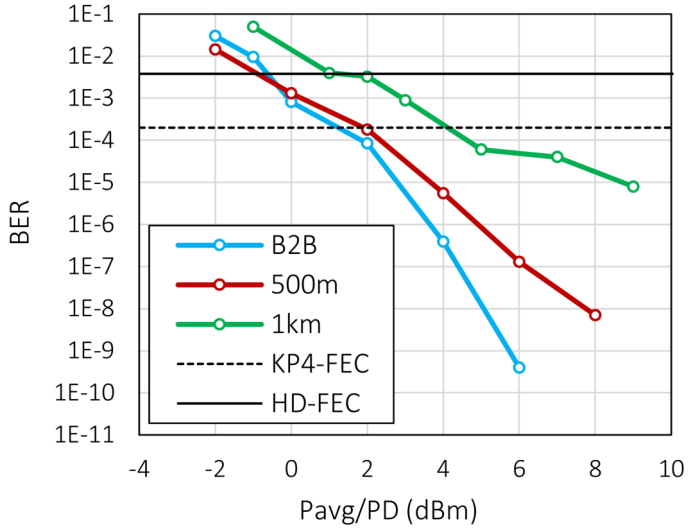

Figure 5.7: Real-time BER curves at 64 Gbaud for back-to-back, 500 m and 1 km of SSMF.

simplify the yield optimization of such a silicon-based transceiver, paving the way towards high-yield and high-volume production.

5.4.2 A path towards O-band EAMs on SOI

Silicon photonics is a very promising platform for data center interconnects, due to its ability to realize compact and low-cost transceivers in high volume leveraging existing CMOS fabrication infrastructure. As demonstrated in this chapter, SiP is capable of delivering very compact transceivers that can operated without the need for DSP, DACs or high-power electronics. The next step would be to combine this modulator with in-house developed dedicated drivers [15] and TIA [19] and realize a low power analog frontend of a silicon transceiver with a dynamic power consumption of less than 2.5 pJ/bit (excluding the laser): i.e. 61 mW to drive the MSB EAM, 45 mW for the LSB (as it needs half the swing), 10-20 mW for the heater and 190 mW for the TIA.

The main downside of the silicon-on-insulator (SOI) platform is that it lacks a native means for optical amplification, and as such no light source. However, many possible solutions exist [20]: lasers can be grown epitaxially on the SOI or butt-coupled to the PIC during assembly without

having to interfere with the SOI. Recently, transfer printing has gained quite some attention as a promising and -most importantly- cost-effective way to selectively transfer pieces of III-V material (e.g. a laser) to a SOI wafer or PIC. The same technique could be used to replace the GeSi EAMs with III-V-based EAMs to allow operation in O-band, as the GeSi Franz-Keldysh effect devices are intrinsically limited to C- and L-band. Another option would be the use of O-band waveguide-integrated EAMs on silicon based on the quantum-confined Stark effect through multiple quantum well SiGe structures, as are currently being developed [21]. However, the proposed optical DAC topology is not limited to EAMs, in fact any intensity modulator could be used.

5.5 Conclusion

We have presented a compact silicon-based transmitter capable of generating 64 Gbaud PAM-4 using two binary driven 120 µm long GeSi EAMs in parallel. Combined with an in-house developed electrical transceiver chipset, we were able to demonstrate the first real-time 64 Gbaud PAM-4 transmission over more than 1 km of SSMF in a chip-to-chip link, without requiring any power-hungry electrical ADCs, DACs or DSP. Integration with the custom-designed BiCMOS drivers of [15] would allow us to realize a 1 pJ/bit transmitter frontend (excluding the laser). These results not only illustrate the advantages of carrying out the DAC operation to the optical domain and thus eliminating the need for linear electronics and optics, but also the capabilities of silicon photonics towards realizing extremely compact and low-power transceivers for 100 Gb/s/λ optical interconnects.

References

[1] J. Verbist, J. Lambrecht, M. Verplaetse, S. A. Srinivasan, P. De Heyn, T. De Keulenaer, *et al.*, "Real-time and dsp-free 128 gb/s pam-4 link using a binary driven silicon photonic transmitter", *Journal of Lightwave Technology*, vol. 37, no. 2, pp. 274–280, 2019.

[2] *IEEE P802.3bs 400 Gigabit Ethernet Task Force.* [Online]. Available: http://www.ieee802.org/3/bs/.

[3] *Consortium for On-Board Optics (COBO).* [Online]. Available: https://onboardoptics.org.

[4] M. Chagnon, M. Osman, M. Poulin, C. Latrasse, J.-F. Gagné, Y. Painchaud, *et al.*, "Experimental study of 112 Gb/s short reach transmission employing PAM formats and SiP intensity modulator at 1.3 μm", *Opt. Express*, vol. 22, no. 17, pp. 21 018–21 036, Aug. 2014. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-22-17-21018.

[5] M. A. Mestre, H. Mardoyan, C. Caillaud, R. Rios-Müller, J. Renaudier, P. Jennevé, *et al.*, "Compact inp-based dfb-eam enabling pam-4 112 gb/s transmission over 2 km", *Journal of Lightwave Technology*, vol. 34, no. 7, pp. 1572–1578, Apr. 2016.

[6] J. Lavrencik, V. A. Thomas, S. Varughese, and S. E. Ralph, "Dsp-enabled 100 gb/s pam-4 vcsel mmf links", *Journal of Lightwave Technology*, vol. 35, no. 15, pp. 3189–3196, Aug. 2017.

[7] W. Wang, P. Zhao, Z. Zhang, H. Li, D. Zang, N. Zhu, *et al.*, "First demonstration of 112 gb/s pam-4 amplifier-free transmission over a record reach of 40 km using 1.3 Directly Modulated Laser", in *2018 Optical Fiber Communications Conference and Exposition (OFC)*, Mar. 2018, pp. 1–3.

[8] A. Samani, D. Patel, M. Chagnon, E. El-Fiky, R. Li, M. Jacques, *et al.*, "Experimental parametric study of 128 Gb/s PAM-4 transmission system using a multi-electrode silicon photonic Mach Zehnder modulator", *Opt. Express*, vol. 25, no. 12, pp. 13 252–13 262, Jun. 2017. [Online]. Available: `http://www.opticsexpress.org/abstract.cfm?URI=oe-25-12-13252`.

[9] J. Sun, R. Kumar, M. Sakib, J. B. Driscoll, H. Jayatilleka, and H. Rong, "A 128 Gb/s PAM4 Silicon Microring Modulator With Integrated Thermo-Optic Resonance Tuning", *Journal of Lightwave Technology*, vol. 37, no. 1, pp. 110–115, Jan. 2019.

[10] J. Lee, S. Shahramian, N. Kaneda, Y. Baeyens, J. Sinsky, L. Buhl, *et al.*, "Demonstration of 112-gbit/s optical transmission using 56gbaud pam-4 driver and clock-and-data recovery ics", in *2015 European Conference on Optical Communication (ECOC)*, Sep. 2015, pp. 1–3.

[11] A. Chiuchiarelli, R. Gandhi, S. M. Rossi, L. H. H. Carvalho, F. Caggioni, J. C. R. F. Oliveira, *et al.*, "Single wavelength 100g real-time transmission for high-speed data center communications", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[12] J. Verbist, J. Lambrecht, M. Verplaetse, J. Van Kerrebrouck, A. Srinivasan, P. De Heyn, *et al.*, "DAC-Less and DSP-Free 112 Gb/s PAM-4 Transmitter Using Two Parallel Electroabsorption Modulators", *Journal of Lightwave Technology*, vol. 36, no. 5, pp. 1281–1286, Mar. 2018.

[13] M. Pantouvaki, S. A. Srinivasan, Y. Ban, P. De Heyn, P. Verheyen, G. Lepage, *et al.*, "Active components for 50 gb/s nrz-ook optical interconnects in a silicon photonics platform", *Journal of Lightwave Technology*, vol. 35, no. 4, pp. 631–638, Feb. 2017.

[14] S. A. Srinivasan, P. Verheyen, R. Loo, I. De Wolf, M. Pantouvaki, G. Lepage, *et al.*, "50gb/s c-band gesi waveguide electro-absorption modulator", in *2016 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2016, pp. 1–3.

[15] H. Ramon, J. Lambrecht, J. Verbist, M. Vanhoecke, S. A. Srinivasan, P. De Heyn, *et al.*, "70 gb/s 0.87 pj/bit gesi eam driver in 55 nm sige bicmos", in *2018 European Conference on Optical Communication (ECOC)*, Sep. 2018, pp. 1–3.

[16] J. Verbist, M. Verplaetse, S. A. Srinivasan, J. Van Kerrebrouck, P. De Heyn, P. Absil, *et al.*, "Real-time 100 gb/s nrz and edb transmission with a gesi electroabsorption modulator for short-reach optical interconnects", *Journal of Lightwave Technology*, vol. 36, no. 1, pp. 90–96, Jan. 2018.

[17] W. Huang, C. Wei, and J. Chen, "Optical dac for generation of pam4 using parallel electro-absorption modulators", in *ECOC 2016; 42nd European Conference on Optical Communication*, Sep. 2016, pp. 1–3.

[18] J. Verbist, M. Verplaetse, J. Lambrecht, S. A. Srivinasan, P. De Heyn, T. De Keulenaer, *et al.*, "100 gb/s dac-less and dsp-free transmitters using gesi eams for short-reach optical interconnects", in *2018 Optical Fiber Communications Conference and Exposition (OFC)*, Mar. 2018, pp. 1–3.

[19] J. Lambrecht, H. Ramon, B. Moeneclaey, J. Verbist, P. Ossieur, P. De Heyn, *et al.*, "56-gb/s silicon optical receiver using a low-noise fully-differential transimpedance amplifier in sige sicmos", in *2018 European Conference on Optical Communication (ECOC)*, Sep. 2018, pp. 1–3.

[20] Z. Wang, A. Abbasi, U. Dave, A. De Groote, S. Kumari, B. Kunert, *et al.*, "Novel Light Source Integration Approaches for Silicon Photonics", *Laser & Photonics Reviews*, vol. 11, no. 4, p. 1 700 063, 2017. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/lpor.201700063.

[21] S. A. Srinivasan, C. Porret, E. Vissers, P. Geiregat, D. V. Thourhout, R. Loo, *et al.*, "High-contrast quantum-confined stark effect in ge/sige quantum well stacks on si with ultra-thin buffer layers", in *CLEO Pacific Rim Conference 2018*, Optical Society of America, 2018, Th3C.1. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=CLEOPR-2018-Th3C.1.

6

4:1 Silicon Photonic Serializer for Data Center Interconnects Demonstrating 104 Gbaud OOK and PAM4 Transmission

With next-generation optical interconnects for data centers aiming for 0.8 Tb/s or 1.6 Tb/s, 100 Gbaud capable transmitters from a single-laser source will become indispensable. However, these lane rates would require bandwidths of 65 GHz or more, doubling the bandwidth requirements of the electrical and optical components with respect to the fastest current generation of optical interconnects running at 53 Gbaud PAM-4.

In this chapter, we propose an integrated 4:1 optical serializer topology to achieve 104 Gbaud NRZ and PAM-4 transmission using only quarter rate components at the transmitter. In a joint experiment with DTU we show that a implementation of this serializer with four GeSi EAMs is capable of delivering 104 (208) Gb/s OOK (PAM4) over 1 km of SMF. For 104 Gbaud OOK, clearly open eyes are obtained, while for PAM-4 the performance is limited by the non-linear E/O-transfer function of the EAM. However, adding pre-emphasis in the electrical driver or replacing the single EAM with the optical DAC topology from the previous chapter, consisting of two EAMs in parallel

with a 90° phase difference between each, could substantially improve these results. Additionally, we discuss the possibility of a four channel transmitter (4 x 208 Gb/s) from a single mode locked laser, amounting to a 832 Gb/s rate based on the current demonstrator.

This chapter is based on the invited journal paper with the same title that was published in the Journal of Lightwave Technology (March 2019) [1].

6.1 Introduction

As the bandwidth requirements for data center interconnects (DCI) keep increasing, the current generation of 100 Gb/s (4 lanes of 25 Gb/s NRZ) single-mode fiber links covering the 500 m to 2 km spans will need to be upgraded to (200 Gb/s or) 400 Gb/s using either (4)8 lanes of 25 Gbaud PAM-4 or (2)4 lanes of 50 Gbaud PAM-4 as described by the IEEE 400GBASE standards. For these 100 Gb/s/λ applications, we have recently demonstrated the first real-time 128 Gb/s PAM-4 link using a binary-driven, silicon-based transmitter without the need of any DSP [2]. However, with the first generation of transceivers operating at 100 Gb/s/λ finding their way to the market in the coming years, research and discussion have already started to shift towards the implementation of the next-generation of optical transceivers operating at 800 Gb/s or even 1.6 Tb/s [3]. To maintain a low channel count, increasing the rates to 200 Gb/s per lane schemes would be an elegant solution for these 0.8 Tb/s and 1.6 Tb/s links.

One option would be to double the amount of bits per symbol again from PAM-4 (2 bit/symbol) to PAM-16 (4 bit/symbol). However, the transition of the NRZ to PAM-4 has already set significant additional requirements on many of the components in the E/O link such as linearity and higher sensitivity (minimally 4.8 dB due to modulation loss and higher when accounting for any additional penalties due to inter-symbol interference [4]), as well as a need for fast DACs and ADCs. All these specifications have lead to substantial increase in the power budget when compared to the simple OOK schemes. Moving to PAM-16 would burden the optical link budget even further, demanding an additional improvement of at least 7 dB in receiver sensitivity, as well as a corresponding increase in the effective number of bits (ENOB) of the DACs/ADCs in the link, likely increasing their already large power consumption much further. Therefore, the prospect of implementing a PAM-16 link seems questionable.

A second path to increase the spectral efficiency of links would be to abandon intensity-modulated direct-detection links (IMDD) and shift to optical coherent technology. The additional signal space, by encompassing the quadrature component, and the increased receiver sensitivity, by using a local oscillator, allows doubling the number of bits/symbol to 4 by employing a 16-QAM modulation format. Yet, data center operators have been reluctant to adopt these metro and long-haul devices as their significantly higher power consumption makes it challenging to meet the extremely compact form factor requirements. Coherent modules targeted for intra data center interconnects such as the recently established 400ZR standard transmit dual-polarization 16-QAM at nearly 60 Gbaud (slightly higher than the typical 53.125 Gbaud rates inside the data center to compensate for the larger overhead caused by the stronger FEC in coherent links) with a power consumption around 15 W. The additional power consumption of the coherent DSP and the large form factor make IMDD still the preferred solution for data center interconnects.

A last straightforward way to increase the line rate is to double the symbol rate to 100 Gbaud, however the implementation is far more burdensome than the concept. The higher symbol rate will require a significant performance increase of the current generation optical and electrical components, requiring bandwidths of 65 GHz or more. Moreover receiver noise will be higher and jitter tolerance will be lower. The challenge in stepping up the baud rate was revealed by the choice of the standardizing committees to go for PAM-4 modulation over the NRZ format for 400 Gigabit Ethernet DCI.

Despite the challenges, an increase of the symbol rate would be a favored solution for DCI and has been subject of much recent research. A 204 Gb/s On-Off keying (OOK) transmitter has been demonstrated using InP platforms for both the optics and electronics in combination with offline DSP [5, 6]. Similarly, an EAM integrated with a DFB laser (EADFB) was used to demonstrate 214 Gb/s PAM-4 transmission [7]. On silicon, the highest reported intensity modulated transmission has been limited to 168 Gb/s PAM-4 using a large multi-electrode traveling-wave Mach-Zehnder modulator (MZM) driven with voltages up to $5\,V_{pp}$ in combination with extensive Tx-side DSP [8], making it less suitable for use in data centers. In line with our previous work, we move functionality to the optical domain to leverage the performance benefits [2, 9] and we focus on a silicon photonic transmitter including a 4:1 serializer to overcome the barriers in stepping up the symbol rates in DCI applications.

Over the last two decades, optical serializers (used in optical time division (OTDM) schemes) have attracted much attention as they can generate very high data rate optical transmissions with limited bandwidth electronics [10]. However the need for long integrated optical delays as well as the absence of a practical and low-cost integrated pulsed laser source has prevented the adoption of optical serializers in datacom applications. Nevertheless, remarkable progress has been made in recent years on integrated semiconductor mode-locked lasers (MLLs) [11, 12]. In particular, the arrival of efficient and low-cost III-V-on-Si MLLs [11] could be an important turning point towards fully integrated Si-based optically serialized transceivers. Silicon Photonics (SiP) would be an ideal platform to integrate MLLs and optical serializers as both devices require long optical time delays, which can be made very compact and with low losses on this platform. Combining a III-V-on-Si laser with the Si multiplexer and a Si modulator would be a promising and cost effective candidate towards realizing 0.8 and 1.6 TbE links, maintaining the low complexity and minimal DSP associated with IMDD links without straining the electronics at the transmitter side.

In this chapter, we demonstrate the first silicon modulator capable of generating 104 Gbaud OOK and PAM-4, using four GeSi electro-absorption modulators (EAMs) as an optical serializer. Driving the EAMs with PAM-4, we achieve the highest reported single-wavelength bitrate for a silicon modulator (208 Gb/s) in an IMDD link. The compact EAMs (80 µm) can be operated without long traveling-wave electrodes and power-consuming terminations (saving transceiver real estate and power) and require only 1.2 V_{pp} at quarter-rate speeds (26 Gbaud) thanks to the 4:1 optical multiplexing. Further improvement is anticipated when the presented topology is enhanced with our SiPh transmitter demonstrated in Chapter 4.

6.2 4:1 Serializer Architecture and Silicon Photonic Implementation with GeSi EAMs

The topology of the serializer is depicted in Fig. 6.1. The discussion will be done for a 4-to-1 serializer, but an equivalent approach can be taken for a 2-to-1 structure or any other integer number. At the input of the circuit, a pulsed light source is required with a repetition rate equal to a quarter of the target symbol rate and duty cycle of 25%, i.e. we have pulses with the length of the symbol period of the target signal. The incoming light is then split equally over four parallel branches. In each branch there is an

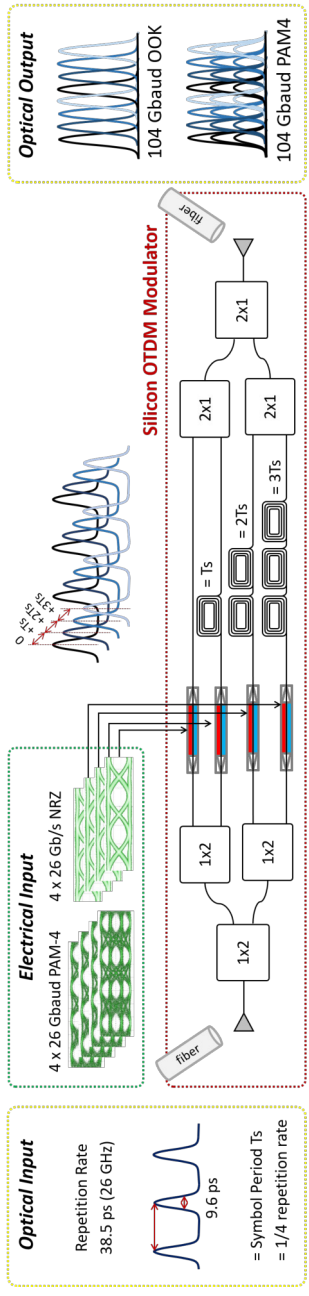

Figure 6.1: Operation principle of an optical 4:1 serializer and modulator: For a target symbol rate Rs, a pulsed light source with a pulse width Ts (corresponding to the desired symbol rate) and a repetition rate at the quarter rate ($Rs/4$) enters the optical chip. The light is divided equally over 4 branches where it is intensity modulated at $fs/4$ and delayed over 0 to 3 symbol periods. After recombination, a waveform is achieved at the target Rs symbol rate.

intensity modulator that is driven by electronics operating at a quarter of the combined target rate. Next, the branches are recombined after applying respective optical time delays of 0, 1, 2 and 3 symbol periods to the modulated pulses. As a result, the pulses are interleaved and form a signal at 4 times the rate of the driving electronics. Moreover, the serializer is operating irrespective of the levels in the pulses, so branches can be modulated with both an NRZ and a multilevel format.

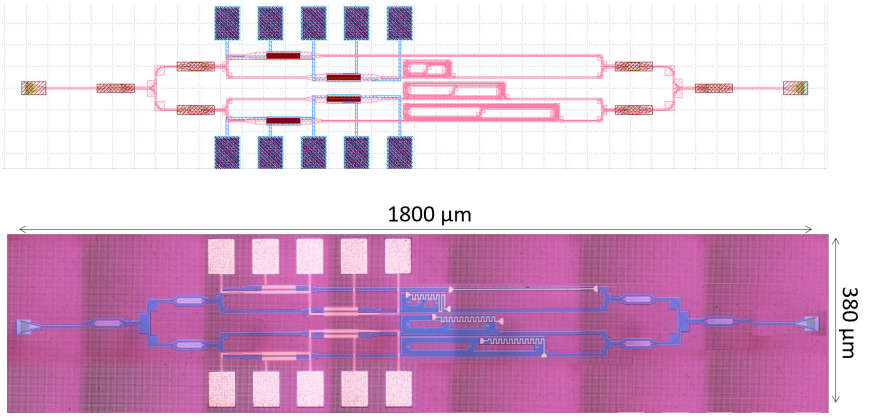

Figure 6.2: Circuit layout and die micrograph.

An implementation of the proposed structure was fabricated on imec's 200 mm SiP platform. The layout and a die micrograph are shown in Fig. 6.2. A binary tree of multimode interferometers (MMIs) is used to split the light into four branches. Per branch, an 80 µm long GeSi EAM modulates the quarter-rate electrical data. The EAM operation is based on the Franz-Keldysh effect, where the bandgap of the device shifts when an electrical field is present. These are the same type of EAMs as used in our previous work [13, 14] where we measured an insertion loss (taken as the ratio of the highest signal power level to the average input power) of ∼7 dB and extinction ratio of ∼8 dB for a 2 V_{pp} swing. As these GeSi devices are extremely compact, they can be driven as very small capacitors without power hungry resistive terminations. Although EAMs were used in this work, any type of modulator structure can be used as long as it implements an intensity modulation of the signal. This opens up the way to also use other compact and unterminated modulators capable of generating PAM-4, such as segmented micro-ring resonators (MRR) [15]. The optical delays on this sample are implemented as waveguide spirals of N×750 µm length (N = 0, 1, 2, 3). Thanks to the high index contrast of the SOI-waveguides, the spirals were realized with a bend radius of 10 µm, resulting in very

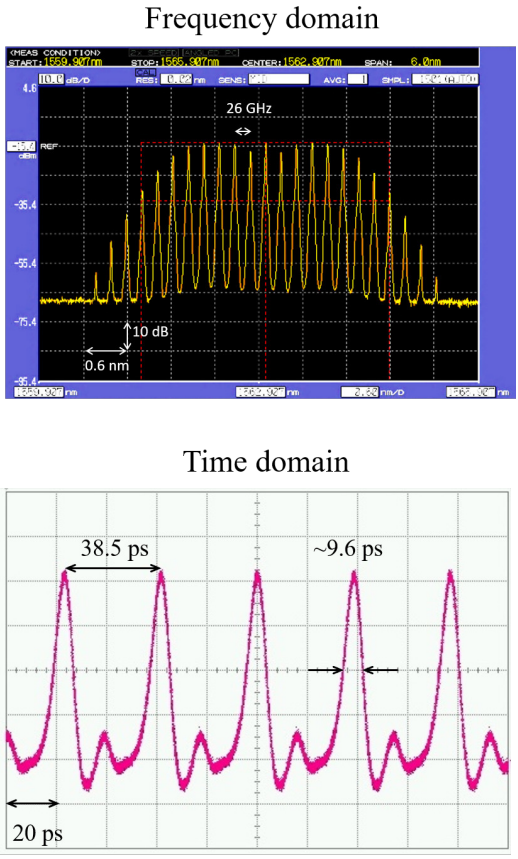

Figure 6.3: Frequency domain and time domain plots of the optical input signal to the optical serializer.

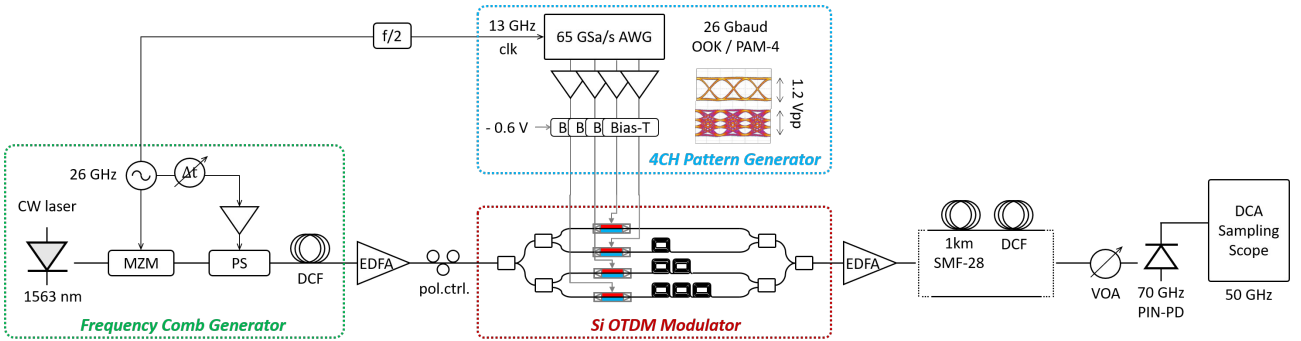

Figure 6.4: Experiment setup for 104 Gb/s OOK and 208 Gb/s PAM-4 transmission with the silicon photonics 4:1 serializer. An MZM pulse carver and a cascaded phase modulator are driven by a 26 GHz sine wave to create a frequency comb from a CW laser. Four channels from the AWG generate data at 26 Gbaud to modulate the pulsed light with NRZ or PAM4 drive signals. The different branches are delayed with spiral waveguides (each 9.6 ps) to properly interleave the modulated pulses into the final optical waveform with a 104 Gbaud symbol rate.

compact time delays with very low losses (<0.2 dB). At 1563 nm, the measured delay difference between each arm is fixed at approximately 9.6 ps. This is at the same time the symbol period Ts of the resulting signal (after serializing), therefore making this demonstrator suited for 104 Gbaud operation originating from 26 Gbaud electronics. Finally, the light from each branch is combined with three 2×1 MMIs and coupled in and out of the photonic IC through fiber-to-chip grating couplers (GC) with an efficiency of approximately -6 dB/coupler. These could be replaced by more efficient edge couplers (<2 dB/coupler) to lower the total insertion loss. A drawback of the configuration, however, remains the inherent 6 dB loss due to combiners after the four branches.

6.3 Transmission Experiment

Fig. 6.4 shows the experimental setup that was used to characterize the fabricated SiP serializer and modulator. To generate a train of optical pulses, CW light at 1563 nm is launched into a cascaded MZM (biased at minimal transmission and acting as a pulse carver) and a phase modulator (PS) modulated by a 26 GHz sine wave generator. 230 m of dispersion compensated fiber (DCF) is added to provide additional compression to the pulses, ensuring a FWHM of less than 9.6 ps when measured with a 70 GHz PIN-PD. Inspection of the optical spectrum reveals eleven 26 GHz spaced lines within a 3 dB flatness (Fig. 6.6). In the time domain signal we can see a pulse periodicity of 38.5 ps, corresponding to a repetition rate of 26 GHz. Next, an EDFA amplifies the pulses to provide a 13 dBm average input power to the PIC. A four channel 65 GSa/s arbitrary waveform generator (AWG) creates four uncorrelated data streams to drive the different branches. These are amplified to 1.2 V_{pp} and applied to the EAMs through two 50 Ω terminated GSGSG RF-probes. All four EAMs were biased at −0.6 V. After the PIC, an EDFA and variable optical attenuator set the average input power to a 70 GHz PIN-PD (0.6 A/W).

Finally, the eyes are captured by a 50 GHz sampling oscilloscope to visualize the eye diagrams. For the BER estimations, a 63 GHz real-time scope is used to capture the waveforms after which the Q-factor is calculated offline. In the experiment, all equipment is synchronized using a 26 GHz clock signal. It is important to have a proper alignment of the pulse carving operation and the AWG data signals to have optimally modulated pulses. When the pulse train generation is replaced by a MLL, active mode-locking would be the natural solution to implement this synchronization. The trigger signal for the MLL and the data signals would

then be retimed using the master clock from the CDR typically present in the electronic driver chip.

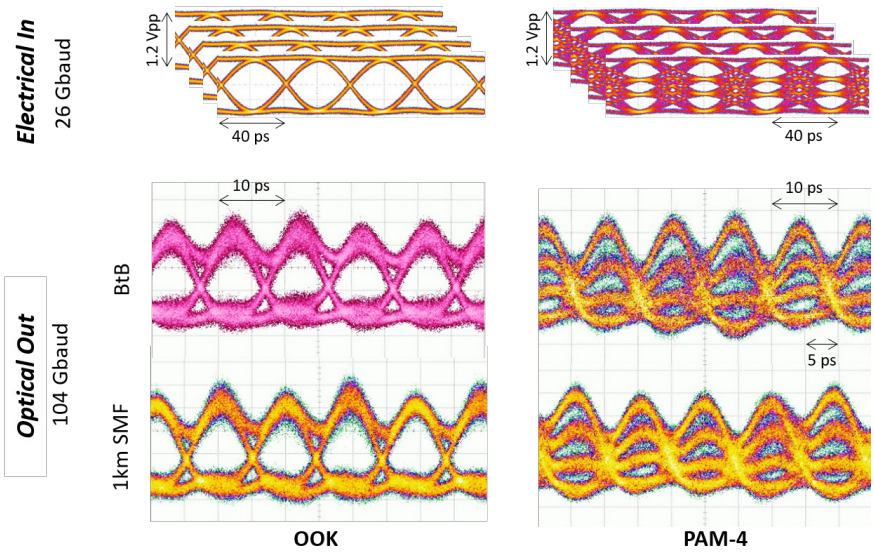

Figure 6.5: Eye diagrams captured by a 50 GHz sampling oscilloscope: 26 Gbaud electrical inputs (OOK/PAM4) and received optical signal in BtB and after 1km of SMF at 104 Gbaud OOK and PAM-4.

The average in-fiber power after the PIC is -14 dBm during operation: 13 dBm (input power) -6 dB (grating coupler in) -7 dB (IL EAMs) -2 dB (7 dB dynamic ER at 1.2 V_{pp}) -6 dB (4:1 combiner losses) -6 dB (grating coupler out). The in-waveguide outer optical modulation amplitude (OMA) after the serializer was −6.8 dBm for an average in-waveguide power of −8 dBm (and +7 dBm in-waveguide power entering the serializer). A considerable part of the loss originates from the 12 dB insertion loss of the grating couplers. Employing edge couplers (<2 dB/coupler) could relieve the total insertion loss and bring the in-fiber power close to -5 dBm. Transmission is tested back-to-back (BtB) and over 1 km of SMF followed by 150 m DCF (-15.9 ps/nm) to compensate for the accumulated chromatic dispersion at 1563 nm. The data transmission was tested with both NRZ and PAM-4 signals (four decorrelated $2^9 - 1$ long pseudo-random bit streams) generated by the AWG at 26 Gbaud.

Fig. 6.5 shows the captured eye diagrams of the input signals to the modulator as well as the resulting 104 Gbaud OOK and PAM-4 eye

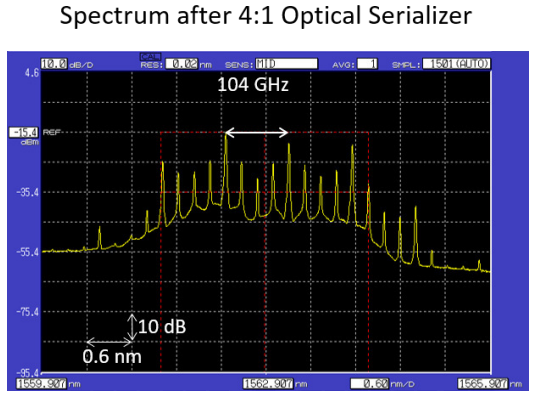

Figure 6.6: optical spectrum after the optical serializer.

diagrams. To capture the individual behavior of each EAM in the eye diagrams, the sampling oscilloscope was triggered with a clock of 1/4th of the total baudrate. However, during BER estimations all eyes were superimposed on one period of the total baudrate (104 Gbaud), resembling a conventional full rate optical receiver. Fig. 6.6 displays the optical spectrum of the output signal. The lines with a 104 GHz spacing are emphasized in the frequency comb revealing the 4:1 serializing operation. At 104 Gb/s OOK, the captured eye diagrams are clearly open, both BtB and after 1 km of fiber. Based on the Q-factor, we estimate a BER below 7E-7 for BtB and 5E-6 after 1 km of SMF, without any online or offline equalization. For PAM-4 transmission, the relative placement of the electrical PAM-4 levels is slightly adjusted in the AWG software to compensate for the compression caused by the non-linear E/O transfer function of the GeSi EAMs. Apart from these adjusted levels, no DSP or equalization is used to generate the 208 Gb/s PAM-4 eyes. Currently, the driving signals for each modulator were not individually optimized. This leads to some unbalancing between the levels of the neighboring tributaries and limits the performance if the signal is received without a demultiplexing operation (as in this experiment). However, a start-up calibration can be implemented to ensure proper alignment, monitoring the optical output generated by a single tributary (placing all other tributaries in a fixed state) and adjusting the driving signal accordingly, and this for all tributaries. Nevertheless, at record speeds of 208 Gb/s, the different PAM-4 levels are still clearly distinguishable even after 1 km of SMF. Here, the estimated BERs are 8.9E-3 for a BtB link and 9.9E-3 after 1 km of SMF, already below the soft decision FEC of 2E−2 (as is commonly used in 200 Gb/s coherent

transceivers) without using any equalization. Adding pre-emphasis and equalization or the DAC-less parallel EAM topology proposed in [14], could be used to improve these BERs for operation with the current standard intra data center FEC for single-mode transceivers of 2.4E-4 for 200 and 400 GbE transceivers.

6.4 Discussion and Prospects

In the demonstrated experiment, we have realized the fastest single-wavelength silicon modulator capable of generating PAM-4 and OOK at 104 Gbaud. The EAM modulators in the different branches are very compact and only need to be driven with 1.2 V_{pp} voltage swing at a quarter of the target symbol rate. As a result, the proposed transmitter topology can be made extremely power efficient. In [16], we have recently shown that for these lumped-load modulators, dedicated drivers can be designed up to 70 Gb/s while only consuming 61 mW. When designed for the 26 Gbaud rate in this demonstrator, the drive current could be lowered, hence resulting in even lower power consumption.

However, there is no fundamental hurdle preventing the proposed architecture to operate at higher data rates. The current photonic IC was operated at 26 Gbaud (quarter-rate) since this was the optimal speed for the measured delays of the waveguide spirals. Only a small redesign effort is required to adapt these delays to shorter symbol periods. Alternatively, when not yet limited by the speed of the electronics, one could choose to implement a half-rate structure rather than a quarter-rate structures to limit interconnect complexity between electronic and photonic dies and gain 3 dB in optical power budget.

A next path forward would be to replace the PAM-4 data generation by an alternative topology. Currently, the four-level signal is generated electrically by the AWG and applied to the non-linear electro-optic transfer function of the EAM. Previously we have shown that by moving the DAC functionality to the optical domain drastically improved PAM-4 performance can be obtained with the same EAMs as it bypasses the linearity requirements at the transmitter [14]. In Fig. 6.7(a) a waveform is displayed where only one of the branches of the serializer is modulated with multilevel data from the AWG, together with the full PAM-4 waveform. We can see that the limitation for the output signal is already in the modulation of a single pulse and that the serializing adds no visible additional degradation to the levels.

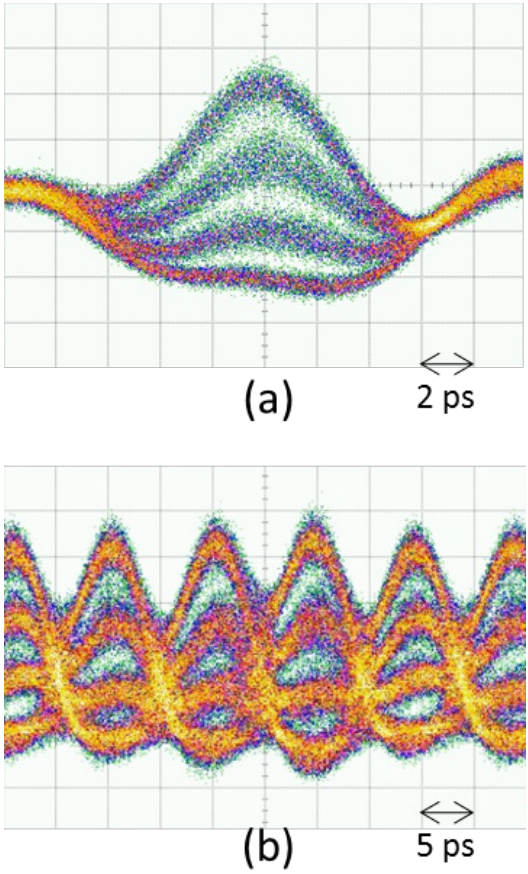

Figure 6.7: Waveform where only 1 out of 4 modulators is driven with four-level data (a) and full PAM-4 waveform (b). The waveforms illustrate that the limitation in the PAM4 signal originates from the modulation of the individual pulses rather than from the serializing operation.

This observation makes it reasonable to assume that employing a better PAM-4 data generation will also result in an improvement when used in the branches of the serializer topology. Ultimately, when combining our 64 Gbaud PAM-4 transmitter demonstrated in Chapter 5 with the serializer, it is possible to evolve towards >0.5 Tb/s from a single laser source. Moreover, the alternative PAM-4 transmitter is again suited to be driven with the same electronic IC as in [16]. This would allow us to achieve these record high data rates with an aggregate driver power consumption of only 1 pJ/bit (excluding the laser).

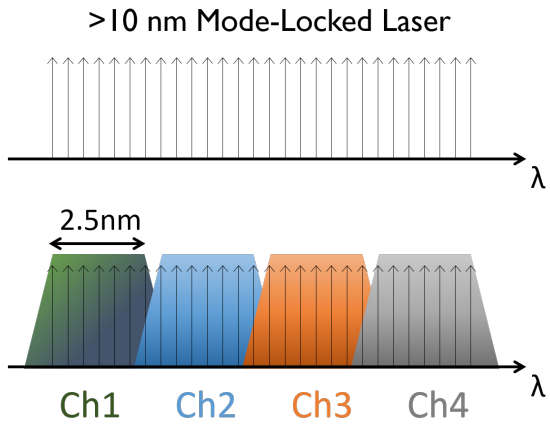

Figure 6.8: Illustrating the principle of using a frequency comb as the source of multiple WDM channels. The minimally required spectral width could be narrowed down to 2.5 nm allowing a single mode-locked laser with a spectrum of at least 10 nm to be shared over 4 channels.

The need for a pulsed laser source remains a disadvantage compared to a common continuous wave laser source as used in today's optical links. To alleviate this drawback, we investigate the possibility of sharing the laser source over multiple channels. Recent research has illustrated that a frequency comb is a viable alternative as a source for multiple wavelength-division multiplexing (WDM) channels [17, 18]. When expanding our test setup with an optical waveshaper acting as a band pass filter, we could narrow down the spectrum of the frequency comb to less than 2.5 nm without any observable penalty. As illustrated in Fig. 6.8, considering a 10 nm frequency comb, a single laser could be used to feed 4 WDM channels. Exploiting this property of the frequency comb can potentially outweigh the additional hardware complexity of the laser source. Replacing the inefficient grating couplers with edge couplers would already provide sufficient power budget to feed four 200G channels from a single mode locked laser (assuming similar output power of the MLL). Additional margin could be obtained by incorporating heterogeneously integrated SOAs either at the input of the transmitter or at the output of each of the four 200G channels. This way, an 832 Gb/s (4×208 Gb/s) PAM-4 transmitter based on a single laser could be realized.

6.5 Conclusion

We have demonstrated a viable path towards increasing the symbol rate in IMDD links by introducing a 4:1 serializer in the optical domain. Using an implementation based on four GeSi EAMs, we show 104 Gbaud OOK and PAM-4 transmission, realizing the fastest PAM-4 transmitter on a silicon platform. In a future implementation, a single 10 nm wide combined with a wavelength demux should be sufficient to feed a 4-channel >800 Gb/s system in a CDWM configuration. The data rates can be pushed even further by omitting the electrical multilevel drive signals and incorporating an optical DAC-based PAM4 generation instead.

References

[1] J. Verbist, M. Vanhoecke, M. Lillieholm, S. A. Srinivasan, P. De Heyn, J. Van Campenhout, *et al.*, "4:1 silicon photonic serializer for data center interconnects demonstrating 104 gbaud ook and pam4 transmission", *Journal of Lightwave Technology*, vol. 37, no. 5, pp. 1498–1503, 2019.

[2] J. Verbist, J. Lambrecht, M. Verplaetse, S. A. Srinivasan, P. De Heyn, T. De Keulenaer, *et al.*, "First Real-Time Demonstration of 128 Gb/s PAM-4 Transmission Over 1 km SMF Using a Si Photonics Transmitter", in *2018 European Conference on Optical Communication (ECOC)*, Sep. 2018, pp. 1–3.

[3] Ethernet Alliance, *Ethernet Roadmap 2018*. [Online]. Available: %7Bwww.ethernetalliance.com%7D.

[4] B. Moeneclaey, J. Verbrugghe, J. Lambrecht, E. Mentovich, P. Bakopoulos, J. Bauwelinck, *et al.*, "Design and experimental verification of a transimpedance amplifier for 64-gb/s pam-4 optical links", *Journal of Lightwave Technology*, vol. 36, no. 2, pp. 195–203, Jan. 2018.

[5] H. Mardoyan, F. Jorge, O. Ozolins, J. M. Estaran, A. Udalcovs, A. Konczykowska, *et al.*, "204-gbaud on-off keying transmitter for inter-data center communications", in *Optical Fiber Communication Conference Postdeadline Papers*, Optical Society of America, 2018, Th4A.4. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=OFC-2018-Th4A.4.

[6] J. M. Estarán, H. Mardoyan, F. Jorge, O. Ozolins, A. Udalcovs, A. Konczykowska, *et al.*, "140/180/204-gbaud ook transceiver for inter- and intra-data center connectivity", *Journal of Lightwave Technology*, vol. 37, no. 1, pp. 178–187, Jan. 2019.

[7] S. Kanazawa, H. Yamazaki, Y. Nakanishi, Y. Ueda, W. Kobayashi, Y. Muramoto, *et al.*, "214-gb/s 4-pam operation of flip-chip interconnection eadfb laser module", *Journal of Lightwave Technology*, vol. 35, no. 3, pp. 418–422, Feb. 2017.

[8] E. El-Fiky, M. Chagnon, M. Sowailem, A. Samani, M. Morsy-Osman, and D. V. Plant, "168-gb/s single carrier pam4 transmission for intra-data center optical interconnects", *IEEE Photonics Technology Letters*, vol. 29, no. 3, pp. 314–317, Feb. 2017.

[9] J. Verbist, J. Lambrecht, M. Verplaetse, J. Van Kerrebrouck, S. A. Srinivasan, P. De Heyn, *et al.*, "Dac-less and dsp-free pam-4 transmitter at 112 gb/s with two parallel gesi electro-absorption modulators", in *2017 European Conference on Optical Communication (ECOC)*, Sep. 2017, pp. 1–3.

[10] T. Richter, E. Palushani, C. Schmidt-Langhorst, R. Ludwig, L. Molle, M. Nolle, *et al.*, "Transmission of single-channel 16-qam data signals at terabaud symbol rates", *Journal of Lightwave Technology*, vol. 30, no. 4, pp. 504–511, 2012.

[11] Z. Wang, K. Van Gasse, V. Moskalenko, S. Latkowski, E. Bente, B. Kuyken, *et al.*, "A iii-v-on-si ultra-dense comb laser", *Light: Science & Applications*, vol. 6, no. 5, e16260–e16260, 2017.

[12] H. Hu, F. Da Ros, M. Pu, F. Ye, K. Ingerslev, E. P. da Silva, *et al.*, "Single-source chip-based frequency comb enabling extreme parallel data transmission", *Nature Photonics*, vol. 12, no. 8, pp. 469–473, 2018.

[13] J. Verbist, M. Verplaetse, S. A. Srivinasan, P. De Heyn, T. De Keulenaer, R. Pierco, *et al.*, "First Real-Time 100-Gb/s NRZ-OOK Transmission over 2 km with a Silicon Photonic Electro-Absorption Modulator", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[14] J. Verbist, J. Lambrecht, M. Verplaetse, J. Van Kerrebrouck, A. Srinivasan, P. De Heyn, *et al.*, "DAC-Less and DSP-Free 112 Gb/s PAM-4 Transmitter Using Two Parallel Electroabsorption Modulators", *Journal of Lightwave Technology*, vol. 36, no. 5, pp. 1281–1286, Mar. 2018.

[15] J. Sun, M. Sakib, J. Driscoll, R. Kumar, H. Jayatilleka, Y. Chetrit, *et al.*, "A 128 gb/s pam4 silicon microring modulator", in *Optical Fiber Communication Conference Postdeadline Papers*, Optical Society of America, 2018, Th4A.7. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=OFC-2018-Th4A.7.

[16] H. Ramon, J. Lambrecht, J. Verbist, M. Vanhoecke, S. A. Srinivasan, P. De Heyn, *et al.*, "70 gb/s 0.87 pj/bit gesi eam driver in 55 nm sige bicmos", in *2018 European Conference on Optical Communication (ECOC)*, Sep. 2018, pp. 1–3.

[17] L. Lundberg, M. Karlsson, A. L. Riesgo, M. Mazur, J. Schröder, and P. A. Andrekson, "Frequency comb-based wdm transmission systems enabling joint signal processing", *Applied Sciences*, vol. 8, p. 718, 2018.

[18] V. Torres-Company, J. Schröder, A. Fülöp, M. Mazur, L. Lundberg, O. B. Helgason, *et al.*, "Laser frequency combs for coherent optical communications", *Journal of Lightwave Technology*, vol. 37, no. 7, pp. 1663–1670, 2019.

7

Coherent Transceivers based on Intensity Modulators

7.1 Coherent communication

Up to this point, the optical links discussed in this work have relied on a single photodiode to detect intensity changes generated by the optical transmitter. However, as illustrated in Fig. 2.2, there are several other physical dimensions of light that can be exploited to carry data. Coherent optical transceivers encode data both in the phase and the amplitude of light, which allows the generation of a higher order constellations such as QPSK and 16-QAM (Fig. 7.1). Besides the phase domain, most coherent links also employ polarization multiplexing as an straightforward way to double the total bitrate as well as the spectral efficiency, i.e. the number of bits per Hz, by utilizing both the TE and the TM polarization states of the light. An illustration of a conventional coherent transmitter with polarization multiplexing that generates 16-QAM, is shown in Fig. 7.2.

Of course, this type of transmitter also requires the conversion of these three physical properties of light (amplitude, phase and polarization) to the electrical domain at the receiver. Whereas in a IMDD link, only the intensity of the optical signal in transduced into an electrical signal

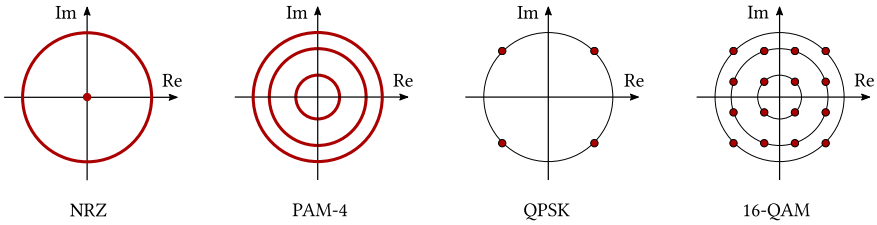

Figure 7.1: Comparison between constellations diagrams of two IMDD formats (NRZ and PAM-4), where the phase can be ambiguous, and two complex modulation formats (QPSK and 16-QAM).

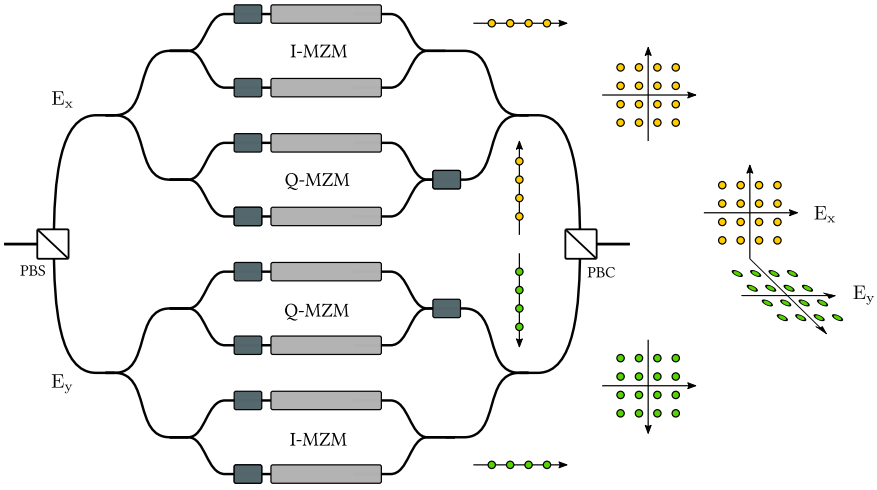

Figure 7.2: Example of a conventional dual-polarization coherent transmitter based on Mach-Zehnder modulators. Each MZM is driven with a four-level signal to produce 16-QAM on both polarizations.

and all phase information is disregarded. This indicates that a coherent optical receiver will have to be much more complex than a simple single photodetector. How this conversion is achieved will be explained more in detailed in Section 7.5.1. The component brake-down of a conventional optical coherent transmitter will be discussed in Section 7.3.

7.2 Coherent Optics for DCIs?

Coherent communication has been studied extensively since the 1980s, when it was developed to increase the optical power budget in long fiber links.

However, with the invention of the EDFA, WDM based systems with in-line amplification negated the need for these highly complex transceivers which at that time still faced many difficulties, and froze the research domain almost 20 years. In the early 2000s, the interest was rekindled as electronic DSPs had become fast enough to track the local oscillator through the received symbols, avoiding the need for an optical phase-lock loop (OPLL) which proved very challenging (and attempts towards a commercially feasible OPLL for coherent telecom applications have only recently started to appear[1]). This evolution made it possible to correct the linear fiber distortions with digital filters, allowing it's processing power to scale down in power and/or up in speed with each new CMOS node according to Moore's law. This also means that these DSPs are typically fabricated in the latest CMOS node, which makes them very expensive.

Because of the success of these digital coherent transceivers, long haul (>100 km) and ultra long haul (>1000 km) links have almost exclusively become coherent-based.

Not more than five years ago, many believed that coherent optical communication would never have a place in the data center ecosystem as they were too big, too power-hungry, and too costly. But as the traffic kept growing at an enormous pace the coherent transceivers began to find their way on many data center road maps (typically somewhere far away in the upper right corner). For this to happen significant changes had to be made to current telecom oriented devices, as DC operators would demand a dedicated solution suited to their specific needs, and not a *light* version of the existing telecom implementations.

One of these areas has been the module size, where remarkable progress has been made in bringing these modules to smaller and smaller packages.

Perhaps surprisingly, in a domain where the performance of the active components is of the utmost importance, it was a silicon photonics based approach that led the way towards more compact and lower power coherent transceivers, for example realizing a complete analog frontend (drivers, TIAs, and PIC) in a ball grid array package operating at 30 Gbaud 16-QAM in 2017[2, 3].

7.2.1 Between data centers: 400G-ZR

Another important step towards bringing coherent to the data center was taken in the spring of 2020, where the first coherent-based multi-vendor implementation agreement for inter-DCI links, 400G-ZR, was adopted by the OIF (Optical Internetworking Forum)[4]. This standard was envisioned to ensure interoperability for coherent inter-DCI modules which can sustain fiber spans of at least 80 km. In 400G-ZR, 16-QAM is sent on two polarizations at nearly 60 Gbaud in C-band rather than O-band due to the lower propagation losses and availability of low-noise and broadband inline fiber amplifiers that can support DWDM. With a maximum power consumption of 15 W housed in a QSFP-DD, OSFP module or on-board optics packaging, a compact and low power transceiver is an important requirement for these application areas.

Therefore, an EAM based solution could be a perfect fit for these links, provided they could generate commonly used coherent constellations such as QPSK and 16-QAM. Building further on the interferometric combination of two EAMs introduced in Chapter 4, Section 7.4 will showcase several possible topologies to do just such a thing.

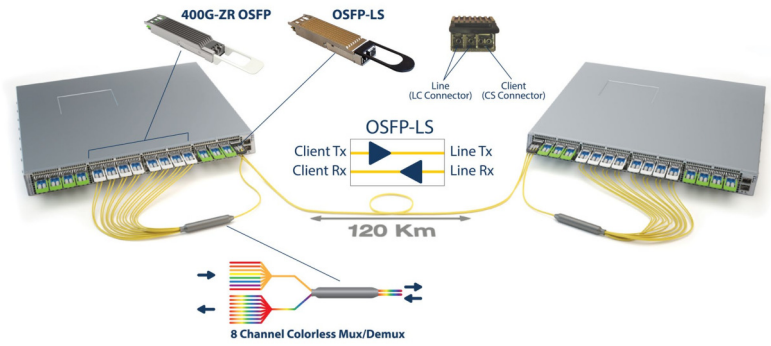

Figure 7.3: Example of a 400G-ZR inter data center interconnect (from Arista [5]).

7.2.2 Inside the data center: coherent versus IMDD

Inside the DC, the power and cost requirements become even more demanding, especially considering the power-hungry and costly DSPs needed for the demodulation of the coherent signal and the compensation of the fiber distortions in a coherent link. However, with the transition from NRZ to PAM-4 in data center interconnects, almost all currently available 400G modules also house a DSP to facilitate, amongst others, the non-linear conversion of data from the electrical to the optical domain and back. To comply with the 400GBASE-DR4 Ethernet standard, the power consumption of these 4×100 Gb/s PAM-4 transceivers is 12 W, which is getting close to the 15 W of the 400G-ZR standard. And as vendors of coherent transceivers like to point out: the differences between their 400G implementation and those of the most challenging links inside the DC is becoming smaller and smaller as can be seen in Fig. 7.4.

This side by side comparison of the blockdiagram of a conventional 400G transceiver module and that of a conventional 16-QAM coherent transceiver for a single polarization, reveals many high-level similarities. Especially, the analog electrical blocks are relatively similar as the drivers and TIAs in both implementations transmit and receive a 4-level electrical signal. Although it must be pointed out that the MZMs in this type of coherent transmitter ideally require twice the swing then if they were to be used in a IMDD link because the modulation ranges from -1 to $+1$ (corresponding to minimum and maximum of the amplitude curves of the MZM transfer function) while in a IMDD link the signals are mapped from 0 to $+1$ (corresponding to the minimum and maximum of power transfer points of the MZM's transfer function).

Yet, there are some significant discrepancies between both versions in other areas of the transceiver. To realize a low phase error on the complex symbols the linewidth of the transmit and the local oscillator laser (i.e. a tapped-off fraction of the transmit laser for the receiving module) needs to remain sufficiently low. Another key difference is found in (necessarily) much stronger DSP capabilities of the coherent chip which adds significantly to the power consumption as well as the latency due to the signal processing and high FEC required for "error-free" transmission, both difficult pills to swallow in a data center environment. Be that as it may, there might be some margin to cut down on the amount of DSP as these chips are built to correct dispersion effects of 100s to 1000s of kilometers. A short reach version up to 2 km might be able strip down significantly in terms of required processing power, levelling the playing

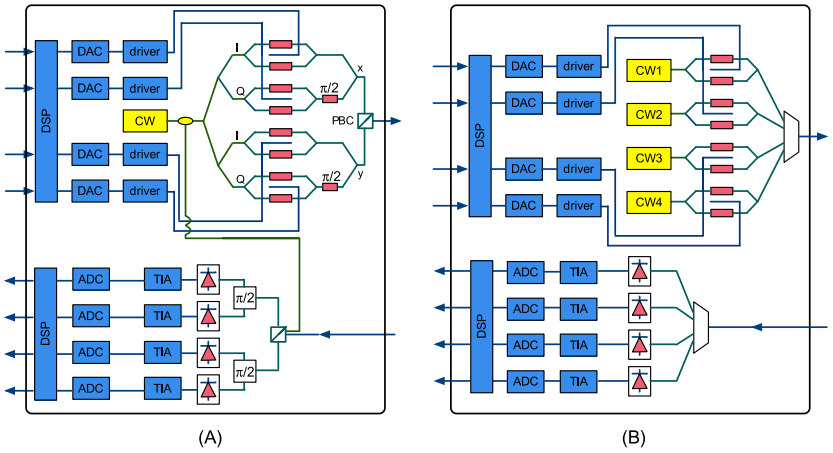

Figure 7.4: Blockdiagram of (a) a typical coherent-based transceiver (here in a dual-polarization version operating at 50 Gbaud 16-QAM per polarization to realize a 400G link) and (b) typical SiPh 400G IMDD-based transceiver (here in a CWDM4 implementation using four wavelengths each carrying 50 Gbaud PAM4 to realize a 400G link).

field somewhat.

Nevertheless, it is still too early to make a confident statement about the chances of coherent optics to have a share of the 0.8 Tb/s or 1.6 Tb/s Ethernet market. It is, however, no longer a question *if* coherent transceivers will have a place inside the data center, but rather *when*.

7.3 A conventional coherent transmitter

Although many different topologies exist to generate complex-valued symbols with both a phase and an amplitude component, by far the most prevalent configuration is the so-called IQ-modulator.

An IQ-modulator consists of two Mach-Zehnder modulators, one for the in-phase modulation (I-MZM) and one for the quadrature modulation (Q-MZM), placed in an interferometer with a fixed 90° phase difference between them. This type of super MZM is referred to as IQ-modulator (or IQ-MZM) as every point of the complex plane can be constructed as the independent superposition of an I and Q signal, i.e. the vector sum of its real and imaginary part.

Each child MZM is biased at its minimum transmission point rather

than its quadrature point, as is the case for IMDD links. This point maximizes the linear modulation region when looking at the fields rather than intensity, as can be seen in Fig. 2.13b. The operating range of the MZM has doubled to be able to generate a full 2π phase shift both in the in-phase and in the quadrature MZM. Consequently, the swing of the electrical drive signal also has to double to be able to reach these phases. Although, due to the cosine transfer function, the conversion of electrical to optical modulation depth becomes much less efficient towards the maximum transmission points of the MZM. With a swing of V_π (half of what is needed for full modulation) already 71% of the total modulation depth is reached. Therefore, the remaining 29% requires the driver swing to increase with an additional V_π.

For constellation with fewer symbols such a QPSK, a reduced optical modulation amplitude might be quite justifiable as the sensitivity penalty at the receiver is likely outweighed by the significant increase in driver power. However, for constellations with a much higher symbol counts (e.g. 16-QAM and higher) one might no longer be able to afford the reduced modulation amplitude. As doubling the driver swing is typically not that easy to do, this means that the already large phase shifters would have to double in length to achieve the same effect, at the cost of increased insertion loss and reduced modulator bandwidth.

In practive, even a swing of "only" V_π is very difficult to obtain, especially on SiP where the $V_\pi L$ values are between 10-20 V mm and even fairly large modulators of 2 mm still would require a swing of 5-10 V_{ppd}. Still significantly higher than a typical (Bi)CMOS driver at 3 V_{ppd} can support. Hence, these modulators tend to be operated with even less modulation depth which increases their insertion loss considerably.

7.4 Coherent Transmitters based on Intensity Modulators

Although the bulk of the power consumption of a coherent module is allocated to the DSP and the laser, there is still a sizeable reduction in power and chip real-estate that can be realized on the transmit side by avoiding conventional IQ-modulator based on large MZMs with relatively low restive terminations (25-30 Ω) which draw larges currents to drive, especially, if the driving signals need to be even stronger than in an IMDD link. Not to mention the possibility of operating at higher baudrates if compact intensity modulators such as EAMs could be used in stead of these

travelling-wave phase shifters. Particularly on silicon platforms, where the weaker plasma dispersion effect tends to lead to higher $V_\pi L$ and thus even longer devices with reduced bandwidths than their InP alternatives. In other words it would be interesting to see if we could also apply the same framework and strategies that were used in chapters 4 and 5 to realize compact, low-cost and low-power IMDD modulators on SiPh, to coherent modulators.

Similar approaches have been demonstrated on InP [6, 7], but have never been attempted on SiPh until very recently [8]. The use-case for this alternative complex modulator on SiPh might even be more compelling than on InP, as SiPh MZMs exhibit a high $V_\pi L$ product combined with more moderate bandwidths than their InP counterparts, requiring either much higher electrical swings (which is already an issue in coherent transmitters), or much longer modulators (which are already several millimeters long and would only reduce the BW as well as increase the insertion losses), or additional signal shaping through analog or digital filtering (both options which will significantly affect the main drawback of the coherent solution: the power consumption). Using a short capacitive modulator such as an EAM or a MRM could potentially save significant transceiver real-estate as well as power, and allow for much higher bandwidths at the same time.

7.4.1 Two-arm QPSK modulator

As QPSK carries the same number of bits per symbol as PAM-4, the PAM-4 DAC topology proposed in Chapter 4 is a natural candidate to be upgraded to QSPK modulator. Recalling the constellation diagram and operating principle of this device (Fig. 4.1 and 4.2), only minor adjustments to produce a 4 symbol complex constellation are required as the topology already uses symbols with different phases to ensure the equidistant position of the PAM-4 levels. It is actually even simpler to generate a square constellation: removing the deliberately introduced OMA difference between both interferometer arms results in a square constellation.

Moreover, the 3 dB power penalty introduced due to the 90° phase difference between both arms, is no longer a drawback of this specific topology as the same penalty also occurs when combining the I and the Q branches of a conventional MZM-based IQ-modulator. This means that total insertion loss for both versions can be very comparable or even lower depending on the insertion loss of the specific stand-alone modulator

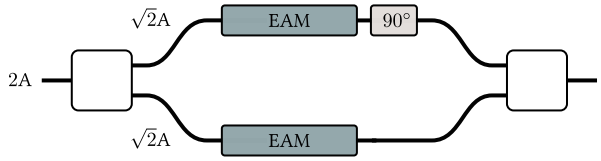

(a) Schematic

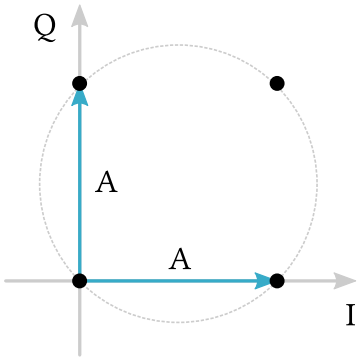

(b) Constellation diagram

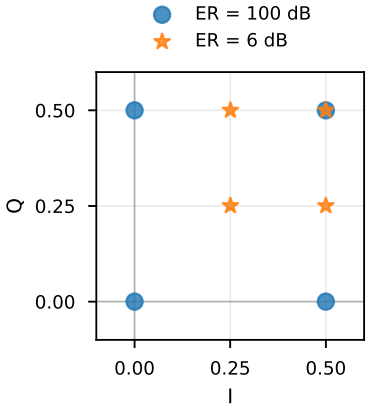

(c) Effect of limited ER

Figure 7.5: QPSK modulator (not DC-centered)

devices, thereby negating a possible advantage of the MZMs over the parallel amplitude modulators. Fig. 7.5 shows such a modulator and its resulting constellation diagram.

Of course, the constellation is now confined in the first quadrant of complex plane which means that it is not DC-centered and that a fraction of the carrier will still be present. Although this is not necessarily a problem, the constellation can be shifted to its ideal position at the cost of slightly higher insertion loss by adding an dummy vector branch with a fixed amplitude $\frac{\sqrt{2}}{2}A$ (with A the shortest distance between two neighboring constellation points which) and phase $-3\pi/4$ as is shown in Fig. 7.6. For EAMs with limited extinction, the amplitude in the dummy branch needs to be increases with a factor $(1+\sqrt{\frac{1}{ER}})$ to compensate for the remaining signal power in the off-state of the EAMs.

There are different options to implement this dummy branch:

- **Star Coupler**

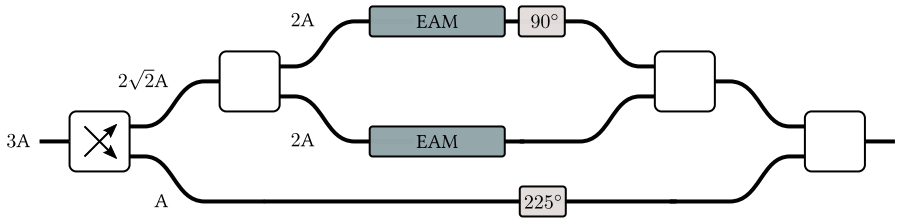

(a) Schematic

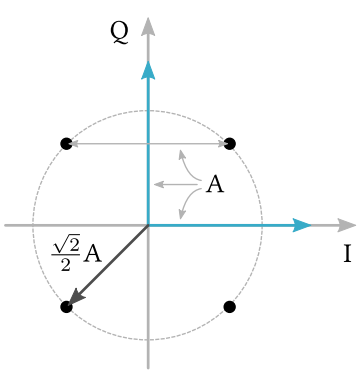

(b) Constellation diagram

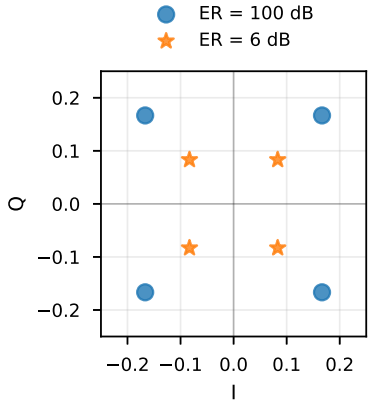

(c) Effect of limited ER

Figure 7.6: QPSK modulator (DC-centered)

A custom-made star coupler could be used to provide a fixed set of amplitude and phase relations for each branch by design. This was the approach followed by [9] in the first demonstration of QPSK modulator using InP-EAMs. Although conceptually elegant, as it requires only two EAMs and two star couplers, there is little tunability left to catch any fabrication deviations. And unfortunately, this was indeed what limited the performance in [9], where the star couplers did not provide the designed phase and amplitude relations.

- **Tunable Power Split**
  Another possibility is to leave out the dummy EAM and work with tunable power splitters which can provide the desired amplitude ratio to the two high-speed arms and the DC-centering arm. This is more efficient, as no power is wasted through the absorption of the dummy EAM to create the desired amplitude ratios. In practice,

tunability will be a desired feature as it allows to compensate for any possible fabrication tolerances during the calibration of the modulator. Therefore, this option will likely be preferred in a commercial product.

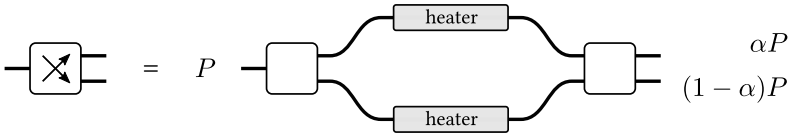

Figure 7.7: Symbol and implementation of a lossless tunable splitter. The incoming power can be divided between both outputs in any arbitrary ratio $\alpha : 1 - \alpha$ with $0 \leq \alpha \leq 1$ by setting one or both heaters to the correct voltage.

This also reveals another interesting property of this modulator: depending on the requirements of the link the modulation can effortlessly be changed between QPSK and PAM-4 by setting the optical power split to the desired ratio (i.e. 50/50 or 33/66).

7.4.2 Four-arm QPSK modulator

Another possibility to generate QPSK with intensity modulators was investigated by [7], where four parallel modulators with a phase of 0°, 90°, 180° and 270° are driven as switches where only one is closed at any given time as illustrated in Fig. 7.8. Although this makes the modulator inherently DC-centered, additional circuitry is required to map every symbol so that only only one of the four modulators is transparent (one-hot switching). The power penalty is 3 dB higher compared to the two-arm EAM-MZI or a IQ-MZM because for each symbol three out of the four arms of the modulator are in a light blocking state.

In stead of driving only one EAM at a time, one could also opt to drive the modulators two-by-two. If we remap the QPSK symbols to 45°, 135°, 225° and 315° and keep the same phase relations between the EAMs, we can see that each symbol can be generated by driving the two neighboring EAMs as can be seen in Fig. 7.9. In Fig. 7.9a every transition to a different symbol takes one charging and one discharging of the modulator capacitance, while in Fig. 7.9b this operation takes 1.33 charges and discharges. However, we gain 41% modulation depth in the optical domain as the amplitude of each symbol is now $\sqrt{2}$ times larger. From Fig. 7.10 we can see that a limited extinction ratio reduces the size of the constellation as can be expected, but does not affect its shape.

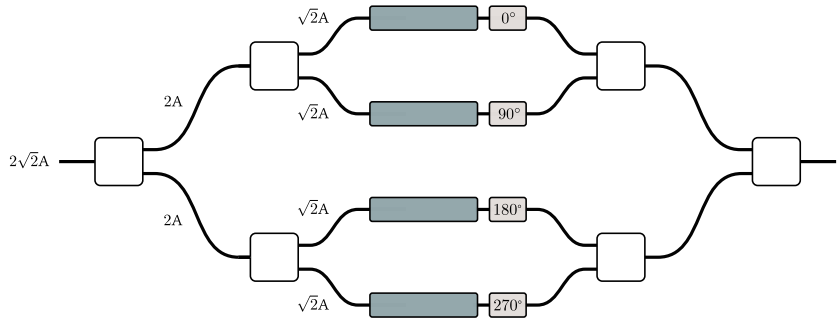

Figure 7.8: QPSK vector modulator with four fixed-phase EAMs.

7.4.3 Push-pull Operation

Looking at Fig. 7.9b, reveals another advantage over the one-hot topology. Each EAM and its 180° rotated counterpart are always driven in a complementary state. This means that we can greatly simplify the bit-to-symbol mapping if we can drive this EAM and its anti-phase twin differentially, i.e. in push-pull. This can be done in two ways:

- **Dual-drive (single bias)**
  Let us call this EAM configuration C-AA-C (A-CC-A), where A and C denominate the anode and cathode of the modulator. To implement this mapping, we only need have to electrically short the cathode (anode) and drive the two anodes (cathodes) with a differential signal. If we assume that both modulators are identical the voltage potential between both modulators becomes an AC-ground and we can arbitrarily assign a DC voltage to it without interfering with the high speed drive signals. This makes this node ideally suited to set the biasing voltage of the modulator as shown in Fig. 7.11a. As this scheme assumes complementary driving signals it matches very well with current mode logic topologies which are often inherently differential. However, a pseudo-differential CMOS topology with invertor-based drivers is also frequently chosen, especially when there are many modulators to be driven as in a segmented modulator [10].

- **Single-drive (dual bias)**
  Another option is depicted in Fig. 7.11b. Here, the modulators are shorted in series in an A-CA-C scheme similar to a balanced photodiode configuration. We might regard this as the dual version of the electro-optical push-pull operation with the AC-CA (or CA-

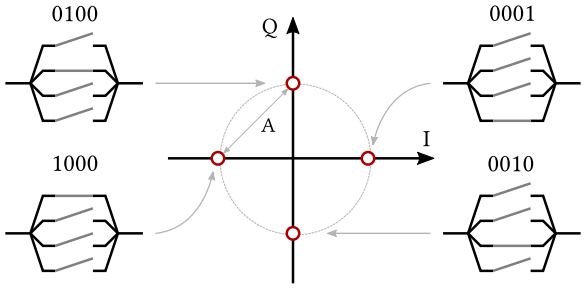

(a) one-hot QPSK

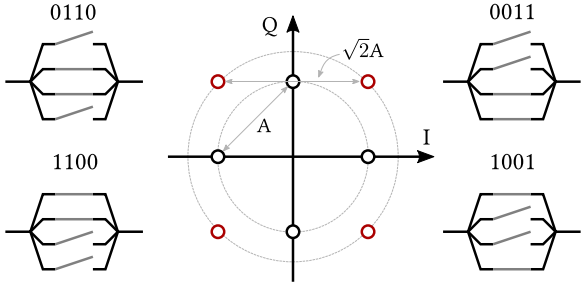

(b) dual-drive QPSK

Figure 7.9: QPSK constellation generated by the quad EAM (a) by driving only one EAM per symbol and (b) by driving two EAMs per symbol.

AC) scheme, where we needed two (differential) drive signals and one biasing voltage. In this configuration, two different bias voltages (e.g. +2 V and 0 V) but only one driver signal is needed to drive both devices push-pull. As the capacitance of these devices is typically very low ($15 - 25$ fF), the disadvantage of higher capacitance is likely outweighed by the reduction in area and complexity as the capacitance of the contacting (bondpads, ESD clamps,...) still dominates the total equivalent capacitive load that needs to be driven. Now, only one -albeit stronger- output stage is needed to drive both EAMs in push-pull. This configuration would be ideally suited for inverter-based CMOS drivers which can provide fast rail-to-rail switching behavior with very high power efficiency.

7.4.4 Binary driven 16-QAM modulator

Higher order constellations are also possible. Just as the PAM-4 or the QPSK modulator was based on a interferometer topology with an

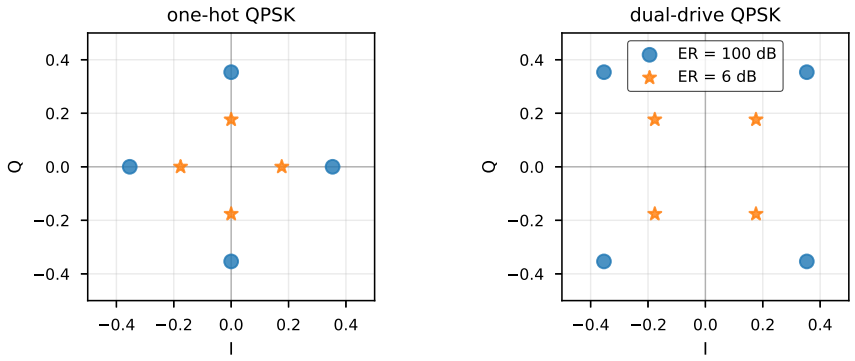

Figure 7.10: Effect of limited ER on a 4-EAM QPSK modulator

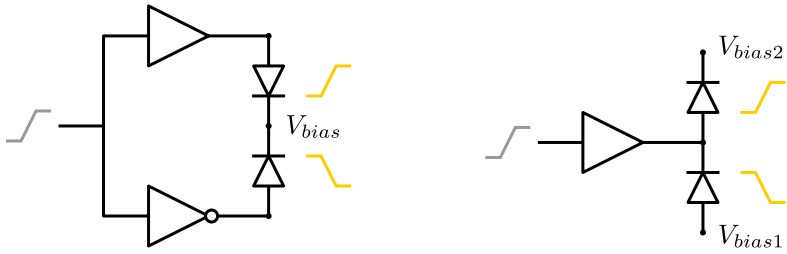

(a) dual-drive push-pull **(b)** single-drive push-pull

Figure 7.11: Two possible ways to bias and drive a two pn-junction-based modulators in push-pull: (a) with two complementary driving signals and one bias voltage ($V_{bias} > V_{cm,drivers}$) and (b) with a single driving signal and two bias voltages ($V_{bias2} > V_{bias1}$)

amplitude modulator in each branch, a 16-QAM modulator can be made by adding two more base (fixed-phase) vectors that make up the modulation space. A four arm interferometer capable of transmitting 16-QAM can be obtained by replacing the EAMs in a PAM-4 vector modulator with the previously discussed QPSK modulator, as can be seen in Fig. 7.12. To achieve a 4-level signal on the I axis with two binary driven EAMs requires a power imbalance between the EAMs corresponding to the 0° and 180° arms (and similarly for the 90° and 270° EAMs on the Q-axis).

As with the two vector version, this constellation is not balanced around DC. However, the same trick can be used as in the 2-arm QSPK modulator to obtain a DC-centered 16-QAM constellation. Adding a dummy branch without high-speed modulation with an amplitude of $\frac{\sqrt{2}}{2}A$ and a phase of

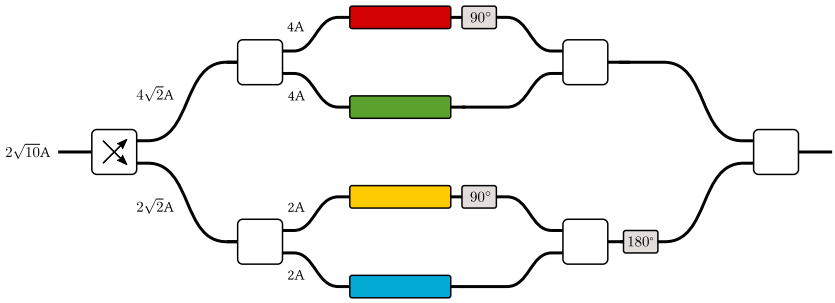

(a) Schematic

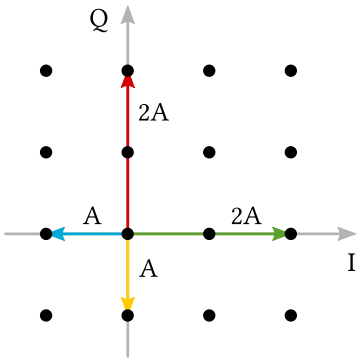

(b) Constellation diagram

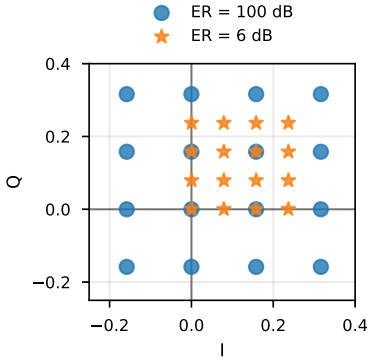

(c) Effect of limited ER

Figure 7.12: Binary driven 16-QAM EAM-based vector modulator (non-DC-centered)

$5\pi/4$ (for ideal EAMs with inifinite extinction) will shift the constellation center back to DC. Furthermore, this solution can be easily adopted to work with more realistic EAMs with limited ER by scaling the amplitude of the dummmy branch with $(1 + \frac{1}{\sqrt{ER}})$, as shown in Fig. 7.13c for an $ER = 100$ dB (close to ideal) and $ER = 6$ dB (more realistic).

7.4.5 Multilevel driven 16-QAM modulator

Until now we have tried to realize coherent transmitters with binary driven modulators to benefit from low-power and CMOS friendly electronics. However, if we abandon this criterium and allow multilevel electrical driving signals, another interesting EAM-based coherent transmitter topology is possible.

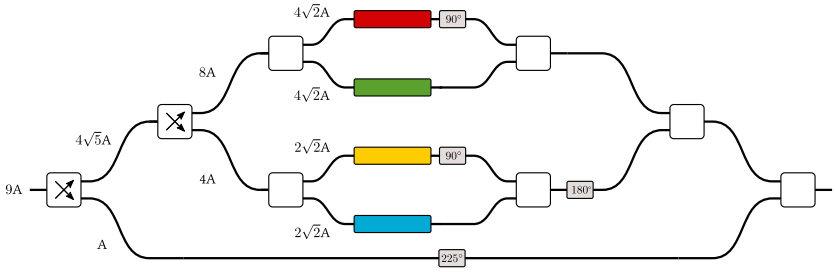

(a) Schematic

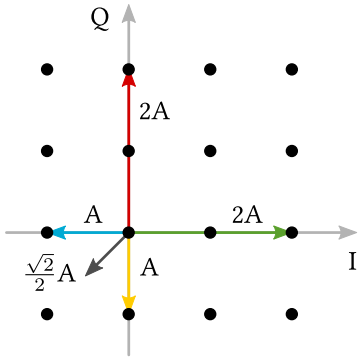

(b) Constellation diagram

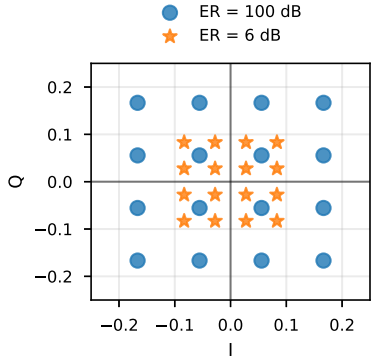

(c) Effect of limited ER

Figure 7.13: Binary driven 16-QAM EAM-based vector modulator (DC-centered)

A disadvantage of the binary driven 16-QAM modulator is that without additional precautions such as a fifth balancing branch, it is inherently unbalanced in DC. Although there was one constellation that did not exhibit this issue: the 4-arm QPSK modulator (Fig. 7.8). Rearranging the prefixed phases so that each child interferometer is put in minimum transmission (180° phase difference) and their modulators driven complementary in push-pull (with two options as discussed in Section 7.4.3), amplitude shift keying can be realized by each anti-phase EAM pair. Fig. 7.14 shows how a 4-level push-pull driver can realize 4-ASK on a single child MZI.

Setting a 90° phase between the inner MZIs, enables true IQ modulation just as it with any conventional MZM-based coherent modulator. Hence, the constellation is inherently symmetrical around DC and this modulator

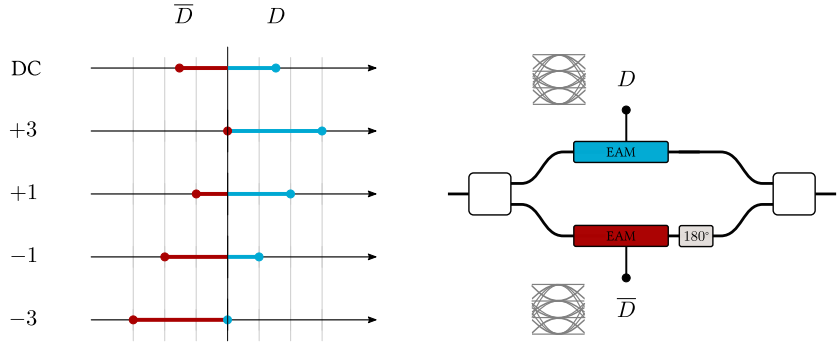

Figure 7.14: 4-ASK operation on two parallel EAMs with PAM-4 in push-pull. The EAM-MZI is biased at minimum transmission.

can therefore act as an in-situ replacement for phase-shifter-based IQ-modulators. And although the four driving signals are no longer binary, only two multilevel differential drivers are needed to operate this EAM-based IQ modulator to obtain 16-QAM modulation.

This means that, in theory, the transmission line drivers from existing travelling-wave based IQ-modulators can be directly reused by simply adding a matching resistive termination. However, significant power reductions can be obtained by being able increase the termination from 25-30 Ω, typical for SiP MZMs, to the conventional 50 Ω as in many other broadband and RF applications. This would reduces the current drawn from the supplies by almost a factor of two, cutting the output driver power almost in half.

However, ideally, the lowest power consumption would come from a close integration of driver and modulator, so that no restive termination is needed and the modulator could be driven as a small lumped capacitor. In [11], a co-designed driver for this type of GeSi EAMs was wirebonded to a single modulator demonstrating up to 70 Gbaud NRZ operation at only 61 mW (excluding the laser power). In a follow-up experiment, a 2-channel version of this driver was used to realize 100 Gb/s PAM-4 with the optical DAC topology from chapters 4 and 5 at only 1.5 pJ/bit.

This EAM-based IQ-modulator strongly resembles the MZM-based IQ-modulator. Both in the electrical domain, where each requires two 4-level differential drivers, as well in the optical domain, where each generates independent amplitude modulations which are combined in quadrature. These similarities allow for a relatively fair side-by-side

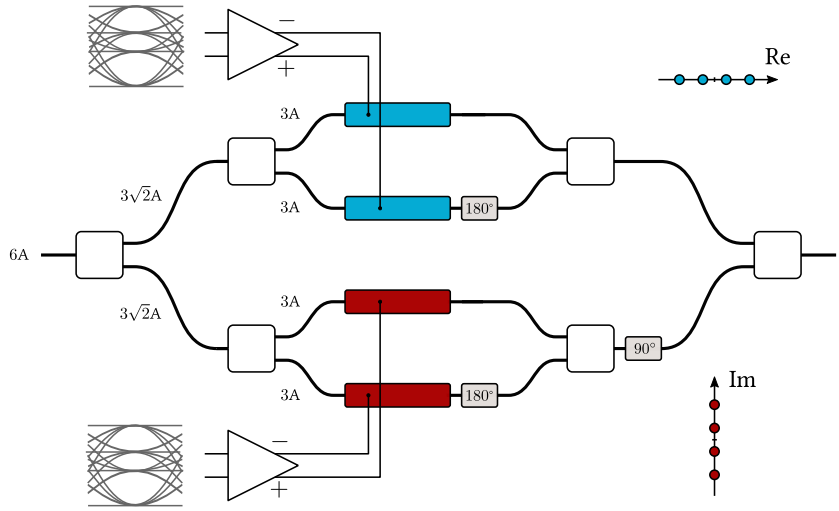

(a) Schematic

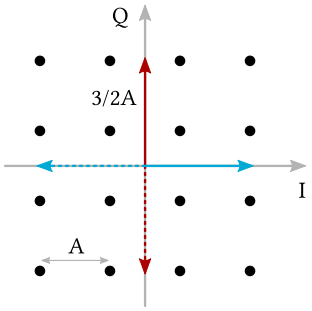

(b) Constellation diagram

Figure 7.15: EAM-based IQ-modulator that can generate DC-centered 16-QAM with two four-level differential driving signals.

comparison between both versions. In Table 7.1 the main properties of both types are listed for 50 Gbaud 16-QAM operation, i.e. a 200 Gb/s on a single polarization.

For the MZM the following assumptions are made: $V_\pi L = 15$V.mm, a length of 1.5 mm to allow a BW close to 35 GHz, and drivers with a differential swing of 3 V_{ppd}.

For the EAMs, the assumptions are based on the 80 µm long SiGe devices from Chapters 3 and 4. With the same 3 V_{ppd} electrical drivers this leads to a extinction ratio around 6 dB.

	PS-based IQM	EAM-based IQM
Electrical input	2×4-level diff.	2×4-level diff.
SE swing per modulator	1.5 V_{pp}	1.5 V_{pp}
driver topology	TML driver (25-30 Ω)	capacitance driver (< 20 fF)
BW	30-35 GHz	>67 GHz
modulator length	1.5 mm ($V_\pi = 10$V)	<0.1 mm
modulation depth	$V_{pp}/V_\pi = 0.15$	ER = 6 dB
IL topology	3 dB	9 dB
IL limited modulation	12.6 dB	6.0 dB
Total theoretical IL	15.6 dB	15.0 dB
IL modulator (EAM/PS)	2 dB	6 dB
Total IL	17.6 dB	21.0 dB
λ operation	O-band or C-band	SiP: C-band, InP: O or C-band

Table 7.1: Comparison between a typical phase shifter IQ modulator and a EAM IQ modulator based on the 80 µm GeSi EAMs from Chapter 3 and 4

In Fig. 7.16, the constellation from both variants are shown. Once with the full modulation and once with limited, but much more realistic modulation depths. Although the EAM-based IQ-modulator has a 6 dB higher theoretical insertion loss to achieve the same constellation, we see that a limited modulation depth is much more penalized on the PS-variant than the EAM-variant. So much so even, that for the assumed modulation depths (MZM: $V_{pp} = 0.15V_\pi$, EAM: $ER = 6$ dB) the EAM-based variant actually needs slightly less input power to obtain a constellation of equal size. However, with an estimated insertion loss of only 2 dB for the PS, the IQ-MZM ultimately has an optical power budget that is 3.4 dB higher.

Nevertheless, the EAM-based IQM is still much more compact and has a BW well over 67 GHz, more than enough BW to support 100 Gbaud operation. The MZM-variant on the other hand will need to be much shorter to achieve a BW around 65 GHz and, hence, will be dominated by an even higher modulation depth penalty. Segmentation of the phase shifters might be a solution to alleviate the BW-limitation, but typically at an even higher electrical power consumption due the high number of segments. Not the mention the much more involved IO between the electrical and the optical domain.

This comparison shows how a coherent transmitter based on EAMs could

be a very attractive alternative to the conventional and bulky MZM-based implementation, especially when targeting baudrates beyond 50 Gbaud.

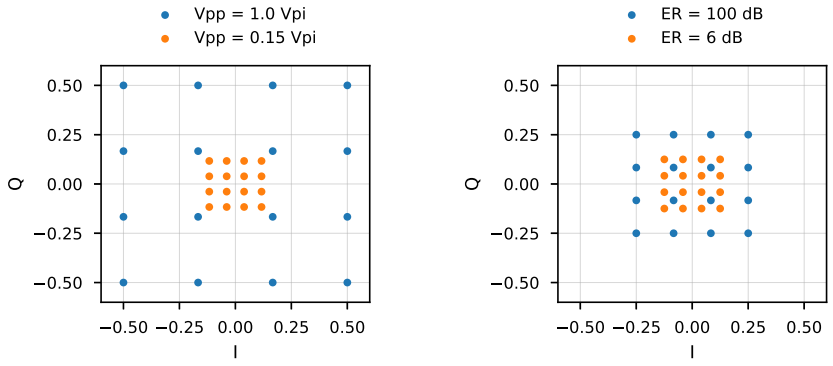

(a) 16-QAM IQ-MZM **(b)** 16-QAM IQ-EAM

Figure 7.16: Comparison of IQ-MZM and IQ-EAM with full modulation and with a 1.5 V_{pp} drive signal per modulator.

7.5 Coherent Receiver

Apart from providing an interesting alternative for a coherent transmitter, EAM-based transmitters also have the distinct advantage that the modulators can be easily set to full absorption and operated as photodiodes. A property that was already exploited in Chapters 3, 4, and 5.

Besides a high responsivity ($\geq$0.8 A/W), the GeSi EAM in particular combines a high BW (>67 GHz) with a flat frequency response. These are especially attractive properties for use in linear receivers as we want as little influence on the O/E conversion as possible. A flat frequency response and good linearity are desirable to minimize the distortion introduced on broadband and non-binary signals such as PAM-4, but these properties are especially important in coherent links.

In these type of links successful reception depends on how well the linear fiber distortions that impeded on the signal can be tracked and compensated, therefore the optical field components have to be converted to the electrical domain with as little non-linear distortion as possible. This places high linearity requirements on coherent links even if the transmitted symbols are not multileveled in amplitude (e.g. BPSK or QPSK).

In this section we will briefly discuss why a coherent receiver is needed and how it functions, and how the four-arm vector modulator can be easily transformed into a single device coherent transceiver.

7.5.1 A conventional coherent Receiver

As a single photodiode acts as a pure power absorber, neglecting any possible phase information, it is not suitable as a receiver in a coherent link. To be able to convert all optical field components (phase and amplitude of each polarization) linearly into the electrical domain, a coherent receiver needs to be much more complex.

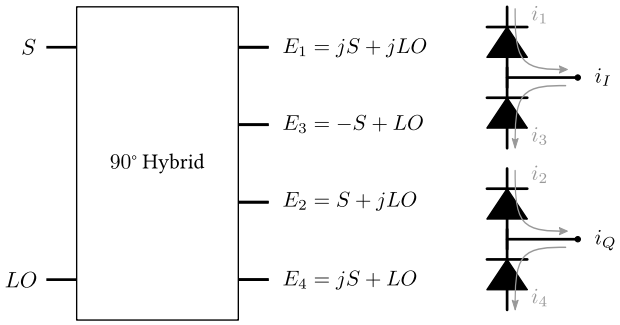

Figure 7.17: Single polarization coherent receiver with 90° degree optical hybrid and balanced photodiodes. (S: received signal, LO: local oscillator)

In a standard coherent receiver as shown in Fig.7.17, the light of an incoming modulated signal is mixed with that of another laser, acting as local oscillator (LO), in a 90° optical hybrid. Typically, only one laser is present in a transceiver module and the LO is a tapped-off portion of the transmit laser. The 90° hybrid is often implemented as a 2×4 multimode interferometers or as a network of four 2×2 multi-mode interferometers. This results in four mixing products (per polarization) each with a 90° different phase relationship between the field of the signal (S) and the LO. In general, the optical fields at the outputs of a 90° hybrid can be written

as:

$$E_1 = \frac{1}{2}(E_s + E_{LO})$$
$$E_2 = \frac{1}{2}(E_s + jE_{LO})$$
$$E_3 = \frac{1}{2}(E_s - E_{LO})$$
$$E_4 = \frac{1}{2}(E_s - jE_{LO})$$

Each mixing product E_i is converted to the electrical domain by a photodiode, yielding the following corresponding photocurrents:

$$i_1 = \frac{1}{4}\mathcal{R}(|E_s|^2 + |E_{LO}|^2 + 2|E_s||E_{LO}|\cos(\Delta\omega t + \Delta\phi))$$
$$i_2 = \frac{1}{4}\mathcal{R}(|E_s|^2 + |E_{LO}|^2 + 2|E_s||E_{LO}|\sin(\Delta\omega t + \Delta\phi))$$
$$i_3 = \frac{1}{4}\mathcal{R}(|E_s|^2 + |E_{LO}|^2 - 2|E_s||E_{LO}|\cos(\Delta\omega t + \Delta\phi))$$
$$i_4 = \frac{1}{4}\mathcal{R}(|E_s|^2 + |E_{LO}|^2 - 2|E_s||E_{LO}|\sin(\Delta\omega t + \Delta\phi))$$

where $\Delta\omega = \omega_s - \omega_{LO}$ and $\Delta\phi = \phi_s - \phi_{LO}$ are the frequency and phase difference between the transmitted signal and the LO, and $\mathcal{R}$ the responsivity of the photodetector. In a homodyne detection scheme where the frequency of the transmitter is equal to that of the LO, $\Delta\omega$ is zero and drops out. It is clear that the amplitude and the phase information of the transmitted symbol is located in the cross product between the transmitted signal and the LO. The first two terms correspond to their respective intensities and are unwanted. As these intensities are identical in all four photocurrents, they can be easily cancelled by subtraction the photocurrents two-by-two. A particularly interesting way to do this is by balancing the photodiodes, as shown in Fig.7.17.

In a balanced detector configuration, the anode and the cathode of the two photodiodes receiving complementary signals (i.e $\{i_1,i_3\}$ and $\{i_2,i_4\}$) are shorted. This current subtraction cancels all the common terms in the mixing products, prevents high DC currents associated with the LO to

enter the TIA, and doubles the sensitivity of the receiver:

$$\Delta i_I = i_1 - i_3 = \mathcal{R}|E_s||E_{LO}| \cos(\Delta\omega t + \Delta\phi) \tag{7.1}$$

$$\Delta i_Q = i_2 - i_4 = \mathcal{R}|E_s||E_{LO}| \sin(\Delta\omega t + \Delta\phi) \tag{7.2}$$

Writing Eq.7.5.1 in terms of powers rather than fields, it becomes clear that next to a phase reference it can also provide a high sensitivity boost:

$$\Delta i_{I/Q} \propto \mathcal{R}\sqrt{P_s P_{LO}} \tag{7.3}$$

The amount of photocurrent corresponds to the geometric mean of the signal power and the LO power, therefore the LO power acts as gain knob on the sensitivity.

As two photodiode currents are merged into one differential current, this configuration also only requires two transimpedance amplifiers per polarization in stead of four, saving cost and power. Another big advantage of the photocurrent subtraction in the balanced configuration, is that the noise related to laser RIN drops out as it is equally present in both photocurrents. On the other hand, balancing the photodiodes almost doubles the capacitive load they present to the TIA, possibly impacting the bandwidth of the optical receiver. This is, however, becoming less and less of a concern as especially in SiPh technologies the capacitance of high-speed PDs has dropped quite a bit over the last 10 years, with state of the art PDs have a junction capacitance of 10-20 fF per PD [12]. Thus -even when doubled due to the balancing- the junction capacitance only represents a minor contribution to the total input capacitance of the transimpedance amplifier, which is typically dominated by the bondpad and the ESD diodes capacitance of the TIA itself.

Fig. 7.18 shows a popular variant of the 90° hybrid using 2×2 MMIs (or any other 2×2 splitter such as a directional coupler) rather than a single-component hybrid such as a 2×4 MMI. The reason is that 2×2 couplers are readily available standard components in any photonic semiconductor fab PDK and as such well characterized, while the 2×4 MMIs typically require custom designs.

In [13], we demonstrated the first silicon coherent receiver capable of receiving 40 Gbaud QPSK and 16-QAM using a 2×4 MMI and a balanced PD configuration together with a linear 2-channel TIA, consuming only 155 mW per channel (excluding the LO laser). Fig. 7.19 shows the demonstrator that was used for these experiments as well as a micrograph of the SiP coherent receiver wirebonded to the 2-channel TIAs.

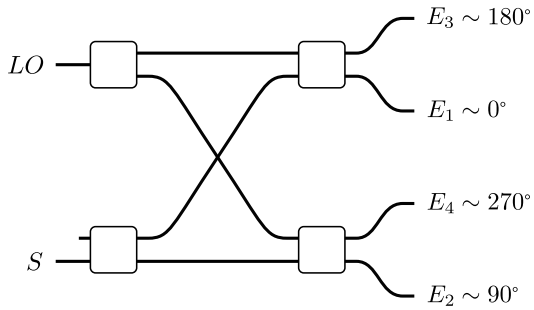

Figure 7.18: Example of a 90° degree optical hybrid implemented as network of 3-dB couplers

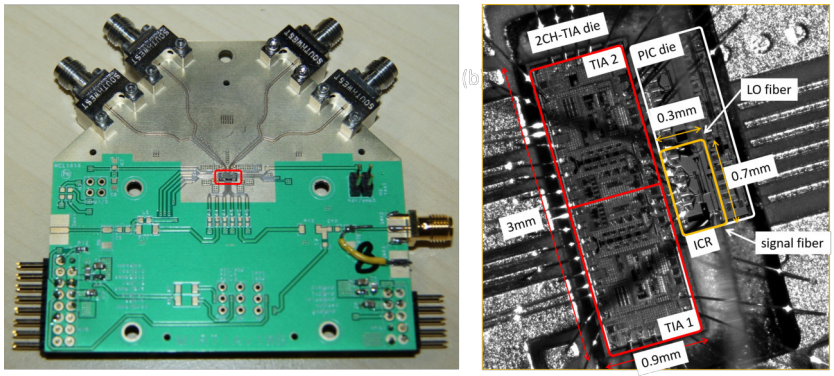

Figure 7.19: Demonstrator for the 40 Gbaud 16-QAM experiments with micrograph of the EIC and PIC from [13].

7.5.2 EAM-based Coherent Transceiver

Although an EAM-based transmitter can be paired with any conventional coherent receiver as described in the previous section, the EAMs themselves could act as an in-situ replacement for the high-speed photodiodes in SiP platforms. For example, in the case of imec's SiP platform, both devices exhibit very similar performance figures in terms of responsivity and bandwidth. This action reduces the number of critical high-speed active components to one, which would benefit the yield optimization efforts on the optical transceiver. Moreover, upon close inspections, it becomes clear that EAM-based coherent receiver shares most of its architecture with a four-arm EAM-based coherent transmitter.

Comparing the intensity modulator based 16-QAM coherent transmitter of Fig. 7.12a and a coherent receiver with a 90° hybrid based on 3-dB couplers, almost all major building blocks are there. We only need to add an additional 3-dB coupler and upgrade the input power splitters of the child interferometers from 1×2 MMIs to 2×2 MMIs. As a consequence, a fixed 90° will now be present by default between both outputs of the 2×2 MMI, which actually takes care of most of the required phase relations in the transmitter for free. Hence, the phase shifters accompanying the EAMs have become optional. Nevertheless, it is prudent to keep them in as they allow to compensate any possible phase imbalances caused by fabrication imperfections of the EAMs and/or the couplers. Fortunately, an additional laser source is not necessarily needed to realize the local oscillator. In many integrated coherent transceivers the laser from the transmitter is reused by dividing its power between the modulator input (transmitter) and the LO (receiver). In this design, the additional optical port of the 2×2 MMI can now serve as input for the incident signal power with (a part of) the Tx-laser acting as the local oscillator, resulting in the device seen in Fig. 7.20a. When operating as a coherent receiver (Fig. c), the thermal phase shifters in the tunable power splitter and in each EAM branch are not used and can set to zero.

Unique to this EAM-based configuration is that it allows half-duplex operation over a single fiber, rather than over two fibers (one for the up link, one for the down link) as is the case in Ethernet based links within and between data centers. Although this implies that there is a convenient way to differentiate which analog electrical circuit is controlling the EAM: the driver or the transimpedance amplifier, as both act on the same device. For RF applications this role can be fulfilled by e.g. a circulator. As the output power of the module laser can be fully used to act as the transmit

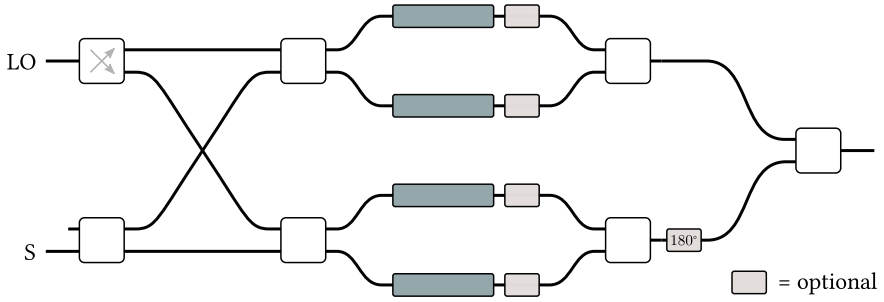

(a) Single-device coherent transceiver

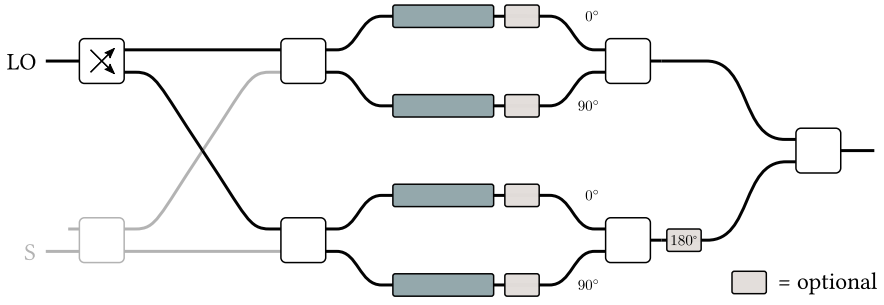

(b) In transmit mode

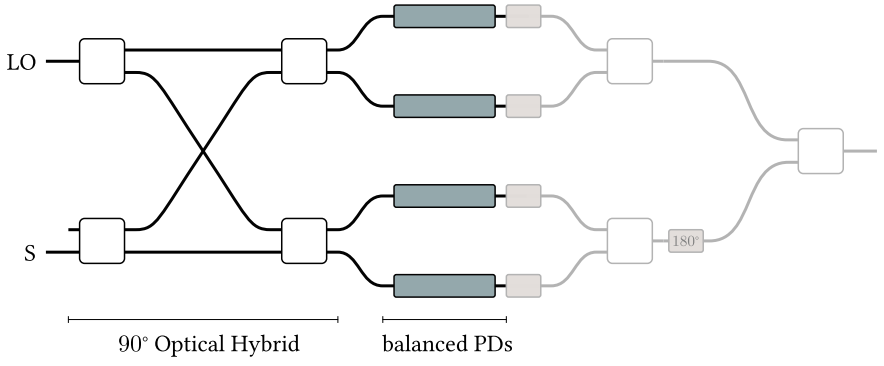

(c) In receive mode

Figure 7.20: Single-circuit coherent transceiver based on EAMs.

carrier in transmit mode or as the LO in receive mode as it does not need to sustain both operations simultaneously, the optical link budget can be several dB higher (higher Tx output power and higher Rx LO power) than that of a conventional coherent transceiver.

Another advantage of being able to reuse the same component for the transmitter and the receiver, is that it greatly simplifies the yield optimization of the optical part of the transceiver: not only does it require a single active high-speed optical device but the transmitter and receiver architecture itself is almost identical.

7.6 Conclusion

In this chapter, we have discussed alternative implementations of coherent transmitters based on EAMs focusing on QPSK and 16-QAM modulators. We described how these intensity modulators can realize complex modulation formats when placed in interferometric topologies, resulting in much more compact designs with lower swings and higher speeds than their MZM based counterparts. This is especially true on silicon where the high $V_\pi L$ of silicon phase shifters, compared to those of InP platforms, lead to large designs and limited bandwidths. On the other hand, EAM-based modulators have a smaller operational wavelength range (a 1 dB link power penalty bandwidth of 30 nm was reported for similar GeSi EAMs in [14]) which should suffice for CWDM implementations but makes them less appealing for dense WDM application, especially when DWDM needs to covers more than one optical transmission bands. An IQ-modulator based on EAMs was discussed, which has a 6 dB higher topological insertion loss than the conventional version with phase shifters when assuming full modulation with ideal modulators. However, with realistic driver swings resulting in limiting modulation depths, the EAM-version becomes more competitive and might even outperform the IQ-MZM in terms of optical power budget as the IQ-MZM suffers from a higher power penalty due to a significantly reduced modulation depth. Lastly, capitalizing maximally on the excellent photodiode capabilities of the EAMs, we have shown that these coherent transmitters can be converted at virtually no additional cost to coherent receivers, realizing a single-circuit coherent transceiver unique to this approach.

Nevertheless, for the foreseeable future, IQ-modulators based on travelling-wave MZMs will remain the go-to topology for long-haul and ultra long-haul coherent fiber links, as they tend to produce higher quality constellations once the formats become much larger than QPSK or 16-QAM.

The EAM-based devices, on the other hand, provide an attractive option to make coherent transceivers the technology of next-generation 800G (e.g. at 100 Gbaud DP-16-QAM) intra and inter data center interconnects, where cost and power consumption are the main drivers in a module's success operating.

Although several devices were designed and fabricated to demonstrate the feasibility and advantages of EAM-based coherent transmitters, unfortunately no time remained to characterize these devices and they will be handed over to a new and eager generation of PhD students. Testimony to the fact that this topic is an active and promising research domain, can be found in the post-deadline paper by Nokia Bell Labs of ECOC 2019 and follow-up journal paper[8] where 100 Gbaud QPSK and 50 Gbaud 16-QAM modulation was demonstrated using the 4-arm differentially-driven GeSi EAMs from Section 7.4.5.

On the receiver side, two experiments were carried out on a (single polarization) integrated coherent receiver based on SiGe photodetectors, a 2×4 MMI as 90° hybrid and a 2-channel linear TIA in 130 nm BiCMOS: basic operation at 28 Gbaud and 40 Gbaud operation in collaboration with Eindhoven university of Technology [13, 15]. At the time of publication, this silicon integrated coherent receiver demonstrated the highest reported baudrate in combination with at a low power consumption of 1.9 pJ/bit (excluding the LO laser) at 40 Gbaud 16-QAM.

References

[1] M. Lu, H.-C. Park, E. Bloch, L. A. Johansson, M. J. Rodwell, and L. A. Coldren, "An integrated heterodyne optical phase-locked loop with record offset locking frequency", in *Optical Fiber Communication Conference*, Optical Society of America, 2014, Tu2H.4. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=OFC-2014-Tu2H.4.

[2] S. Hardy, *Acacia communications samples cfp2-dco with internal coherent dsp*, 2016 (accessed: 01.16.2020). [Online]. Available: https://www.barrons.com/articles/acacia-can-ride-the-rise-of-silicon-photonics-says-mkm-1505777100.

[3] C. Doerr, J. Heanue, L. Chen, R. Aroca, S. Azemati, G. Ali, *et al.*, "Silicon photonics coherent transceiver in a ball-grid array package", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[4] O. P. Release, *Oif publishes implementation agreement for 400zr coherent optical interface.* [Online]. Available: https://www.oiforum.com/oif-publishes-implementation-agreement-for-400zr-coherent-optical-interface/.

[5] Arista, *Data Center Interconnect - Simplified.* [Online]. Available: %7Bhttps://www.arista.com/en/11196-webinar-dci-with-400g-osfp%7D.

[6] C. R. Doerr, L. Zhang, P. J. Winzer, and A. H. Gnauck, "28-gbaud inp square or hexagonal 16-qam modulator", in *Optical Fiber Communication Conference/National Fiber Optic Engineers Conference 2011*, Optical Society of America, 2011, OMU2. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=OFC-2011-OMU2.

[7] G. de Valicourt, H. Mardoyan, M. A. Mestre, P. Jennevé, J. C. Antona, S. Bigo, *et al.*, "Monolithic integrated inp transmitters using switching of prefixed optical phases", *Journal of Lightwave Technology*, vol. 33, no. 3, pp. 663–669, Feb. 2015.

[8] A. Melikyan, N. Kaneda, K. Kim, Y. Baeyens, and P. Dong, "Differential drive i/q modulator based on silicon photonic electro-absorption modulators", *Journal of Lightwave Technology*, vol. 38, no. 11, pp. 2872–2876, Jun. 2020.

[9] C. R. Doerr, L. Zhang, A. L. Adamiecki, N. J. Sauer, J. H. Sinsky, and P. J. Winzer, "Compact eam-based inp dqpsk modulator and demonstration at 80 gb/s", in *Optical Fiber Communication Conference and Exposition and The National Fiber Optic Engineers Conference*, Optical Society of America, 2007, PDP33. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=OFC-2007-PDP33.

[10] M. Vanhoecke, A. Aimone, N. Argyris, S. Dris, R. Vaernewyck, K. Verheyen, *et al.*, "Segmented Optical Transmitter Comprising a CMOS Driver Array and an InP IQ-MZM for Advanced Modulation Formats", *Journal of Lightwave Technology*, vol. 35, no. 4, pp. 862–867, Feb. 2017.

[11] H. Ramon, J. Lambrecht, J. Verbist, M. Vanhoecke, S. A. Srinivasan, P. De Heyn, *et al.*, "70 gb/s 0.87 pj/bit gesi eam driver in 55 nm sige bicmos", in *2018 European Conference on Optical Communication (ECOC)*, Sep. 2018, pp. 1–3.

[12] H. Chen, P. Verheyen, P. D. Heyn, G. Lepage, J. D. Coster, S. Balakrishnan, *et al.*, ''-1 V bias 67 GHz bandwidth Si-contacted germanium waveguide p-i-n photodetector for optical links at 56 Gbps and beyond", *Opt. Express*, vol. 24, no. 5, pp. 4622–4631, Mar. 2016. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-24-5-4622.

[13] J. Verbist, J. Zhang, B. Moeneclaey, W. Soenen, J. Van Weerdenburg, R. Van Uden, *et al.*, "A 40-gbd qpsk/16-qam integrated silicon coherent receiver", *IEEE Photonics Technology Letters*, vol. 28, no. 19, pp. 2070–2073, Oct. 2016.

[14] S. A. Srinivasan, P. Verheyen, R. Loo, I. De Wolf, M. Pantouvaki, G. Lepage, *et al.*, "50gb/s c-band gesi waveguide electro-absorption

modulator", in *2016 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2016, pp. 1–3.

[15] J. Zhang, J. Verbist, B. Moeneclaey, J. Van Weerdenburg, R. Van Uden, H. Chen, *et al.*, "Compact low-power-consumption 28-gbaud qpsk/16-qam integrated silicon photonic/electronic coherent receiver", *IEEE Photonics Journal*, vol. 8, no. 1, pp. 1–10, 2016.

8

Conclusion & Outlook

8.1 Summary of the results

We began this book by taking a step back and looking at the bigger picture to see which trends fuel the research activities in this specific research area, more specifically the almost exponential growth of the internet and the boom of the data center industry in the last ten years. As our consumption and generation of data keep increasing, the need for faster, cheaper and more efficient optical interconnects is at an all time high.

During this research, we have investigated and experimentally demonstrated several ways to realize compact and low-power silicon photonics transceivers, striving for topologies that can be driven with low swing (<2 V_{pp}) binary signals to ensure CMOS compatibility, by using waveguide-integrated GeSi electro-absorption modulators on a 220nm SOI platform.

In Chapter 3 we demonstrated real-time 100 Gb/s transmission using NRZ and electrical duobinary on a single GeSi EAM in combination with a BiCMOS transmitter and receiver chipset. This was the first reported 100 Gb/s NRZ (and EDB) transmission in silicon photonics without any need for digital signal processing. In another experiment, the same EAM on a second die replaced the external IIIV-based photodiode as optical receiver, resulting in the first single-lane silicon photonics based link at

100 Gbaud.

In the following chapter, a novel optical modulator topology is presented that allows the optical generation of PAM-4 based on the vector sum of two binary driven amplitude modulators in parallel. The role of the phase difference and power split between both modulators and their effect on the placement on the PAM-4 eye is discussed in more detail, as well as their possible application into realizing predistorted PAM-4 eyes.

A demonstrator of this optical DAC was fabricated using GeSi EAMs, which clearly outperformed a single multilevel driven modulator in terms of eye quality irrespective of the symbol rate, thereby validating the advantage of an optical DAC topology over an electrical topology in terms of performance. These experiments resulted in the first 100 Gb/s and 112 Gb/s PAM-4 transmission on SiP without requiring any DSP, power-hungry DACs or on-chip terminations.

Chapter 5 continued with the same modulator topology to showcase the first real-time link experiments at 128 Gb/s PAM-4 with a silicon transmitter, as well as a full real-time silicon link where the receiver is a strongly reversed biased EAM. We also showed that an optically induced power difference between the LSb and MSB leads to a higher OMA than a electrically induced difference, while an electrically induced difference can be used to save power as one of the drivers operates at a significant lower swing.

Next, we turned our attention towards next-generation of optical interconnects in Chapter 6. Here, we demonstrated how an optical serializer could be used to leverage the current or even previous generation modulators to achieve beyond 100 Gbaud lane rates for both NRZ and PAM-4, and how this -in combination with a pulsed light source such as mode lock lasers- could be used to realize next-generation transceivers operating at 0.8 Tb/s or more. With this experiment we showed the highest reported IMDD transmission rate on silicon without DSP at 208 Gb/s PAM-4 with an integrated 4:1 optical serializer, requiring only quarter-rate (26 Gbaud) electronics and photonics.

And finally, in Chapter 7, another path towards the future of optical transceivers is explored, where we investigated how the design considerations set forth in the previous chapters can be used to realize compact and low-power coherent transceivers using EAMs. Or in other words, how to mimic phase modulation without using phase modulators, either with binary or with multilevel driving signals.

8.2 Alternative PAM-4 ODACs using EAMs

Double-Sided EML

Performing the DAC operation to generate PAM-4 in the optical domain has had several advantages as shown throughout Chapters 4 and 5. And although SiPh might be more suited towards low-cost implementations, the designs proposed in this book are platform agnostic and could be translated easily to IIIV-based platforms such as InP. One particularly interesting use-case could be found in a double-sided EML.

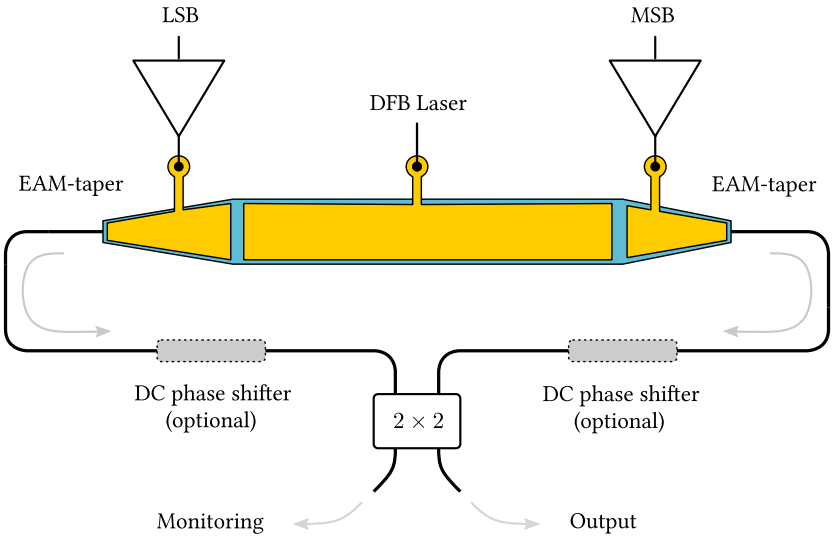

Figure 8.1: A two-side EML or a EAM-DFB-EAM structure where each of the tapers coupling light from the laser gain medium to the waveguide can act as an independent EAM. Combining both waveguides with a 2×2-MMI results in a 90° fixed phase difference between LSB and MSB. As long as the there is a 3 dB difference in OMA between the LSB and MSB, the PAM-4 signal will be perfectly equidistant as described in Chapter 4. The other output provides the complementary signal (where the MSB is in quadrature rather than the LSB) and can be used as a monitoring port. For an optically introduced LSB-MSB power split a DFB with an asymmetrically placed quarter wave shift could be used, resulting in a different output power from each side of the EML.

The optically generated PAM-4 topology introduced in Chapter 4 would be a perfect fit for a two-sided EML, i.e. a EML with electro-absorption section at both sides of the laser as for example shown in [1, 2]. In [3],

we demonstrated such a IIIV-on-Si two-sided EML where the EAMs and the output power at both sides are almost identical, and where we could generate 50 Gb/s NRZ signals and 100 Gb/s duobinary signals from each EAM. Reducing the optical swing from one EAM to half that of the other and combining both waveguides with a 90° phase difference (e.g. standard in a 2×2-coupler), would results in a linear PAM-4 signal. A schematic of such a PAM4-EML is shown in Fig.8.1.

Dual-polarization ODAC

In all previous examples of the PAM-4 ODAC, care had to be taken when interferometrically combining the LSB and MSB signals. As we explained in Chapter 4, it is impossible to produce an equidistant PAM-4 signal without introducing a 90° phase difference between both branches.

Because, when both modulators are in the "1"-state, the sum of the LSB and the MSB signals would result in a mixing product, after reception with a square-law photodetector, which will always lead to a too high power level:

$$\begin{aligned} i_{ph} &= \mathcal{R}|E(t)|^2 \\ &= \mathcal{R}|E_{MSB} + E_{LSB}|^2 \\ &= \mathcal{R}(P_{MSB} + P_{LSB} + 2\sqrt{P_{LSB}P_{MSB}}) \end{aligned}$$

To remedy this unbalance, a 90° phase difference is introduced in the LSB or the MSB branch (here, in the LSB):

$$\begin{aligned} i_{ph} &= \mathcal{R}|E_{MSB} + e^{j\pi/2}E_{LSB}|^2 \\ &= \mathcal{R}|E_{MSB} + jE_{LSB}|^2 \\ &= \mathcal{R}(P_{MSB} + P_{LSB}) \end{aligned}$$

This will cause the mixing product to drop out of the equation, resulting in equidistant power levels irrespective of the extinction ratio, insertion loss, or transfer function of the modulators as long as they are sufficiently similar. Unfortunately, interferometrically combining two signals in quadrature comes at an insertion loss of 3 dB, because we actively avoided full constructive interference in the optical power combiner by adding the 90° phase shift.

Yet, one can circumvent this power penalty due to the interferometric recombination and the square-law behavior of the photodetector by

avoiding the interferometric recombination all together. This can be done by placing the LSB and the MSB optical signals on another polarization, e.g. the MSB on the TE polarization and the MSB on the TM polarization, as is shown in Fig. 8.2. When captured by a (polarization insensitive) photodetector, both polarizations will independently be converted to current without introducing a mixing product:

$$\begin{aligned} i_{ph} &= \mathcal{R}|E_{TE}|^2 + \mathcal{R}|E_{TM}|^2 \\ &= \mathcal{R}(P_{MSB} + P_{LSB}) \end{aligned}$$

Typically, the waveguides and optical components are optimized for TE polarization on chip, so the polarization rotation and combination ideally needs to happen right before the light is launched into the fiber. On a silicon platform all three functionalities can even be achieved simultaneously with a dual-polarization or 2D-grating coupler [4].

A similar approach was taken in [1] using a double-sided InP-based EML, but the polarization handling (rotation and combination) had to be performed off-chip as it is not readily available on InP platforms. Therefore, the previously discussed ODAC might have been a better fit. Alternatively, a IIIV-on-Si double-side EML would allow both variants, interferometric or polarization diverse, to be explored as the polarization handling could be performed in the underlying silicon platform.

To receive both polarization on a PIC-based photodetector one can use a polarization-independent edge coupler and a regular photodetector (Fig. 8.3a), or a dual-polarization grating coupler and feed both outputs to each end of the photodiode (Fig. 8.3b). A power combiner can not be used as this would introduce the interferometric dependence that we are trying to avoid. The first option has the additional advantage that - although it uses both polarization- the reception never depends on the specific polarization state but only on the fact the both polarizations are orthogonal to each other. For the latter option, the polarizations have the be aligned with expected states of the dual-polarization grating coupler. This can be done by using polarization maintaining fiber, however these fibers are considerably more costly than regular SMF and much more complex to in terms of optical alignment. Of course, if the photodetector does not need to come from the same PIC as the transmitter many other other options exist as long as the trajectory from the fiber output to the photodetector is not polarization sensitive. For example, top or bottom illuminated PDs, as commonly used in IIIV-based optical transceivers, are polarization insensitive and therefore can be paired directly with this type of transmitter [5].

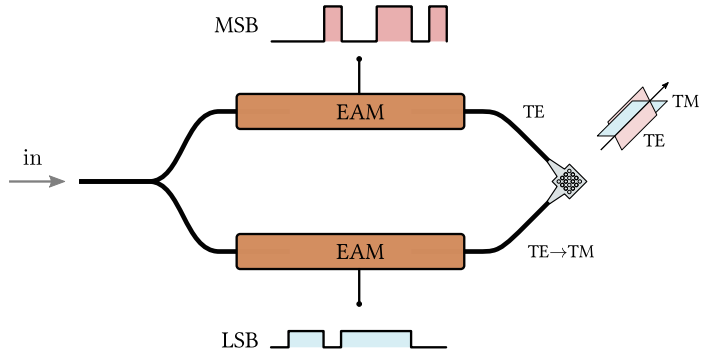

Figure 8.2: PAM-4 ODAC alternative where the LSB and MSB are carried by different polarization states. Polarization combining as well as rotation of the TE mode to TM is achieved by a 2D grating coupler.

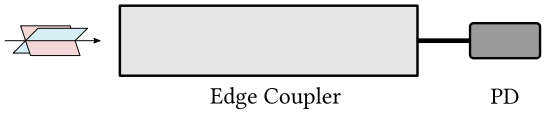

(a) Receiver with polarization independent edge coupler

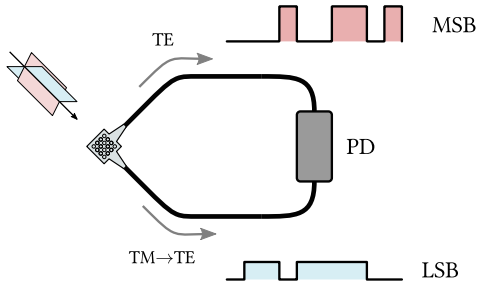

(b) Receiver with polarization dependent 2D-grating coupler

Figure 8.3: Two examples of integrated optical receivers topologies for receiving a dual-polarization PAM-4 signal.

8.3 DCIs: what comes next?

This is the million dollar question. As more and more suppliers of pluggable modules are finalizing their product for 200G/400G Ethernet, the industry starts to look towards what will become the next-generation data center interconnect. Many questions are still to be answered. Will there still be a place for pluggable modules? Or will on-board optics finally take off? What will be the throughput of this new generation interconnects: 0.8 Tb/s or 1.6 Tb/s? Either way these will most-likely require 200G optical lane rates. Will it possible to increase the modulation order further (e.g. 80 Gbaud PAM-6 or 70 Gbaud PAM-8) or will the optics and electronics be able to support 100 Gbaud PAM-4?

Currently, the latest state-of-the-art electronics have demonstrated that it is feasible to generate 100 Gbaud PAM-4 with a relatively low power consumption (700 mW) by linearly interleaving PAM-4 streams[6], similar to the optical serializer in Chapter 6. Receiving this signal optically is feasible with the current generation of photodiodes, as both InP and SiPh have demonstrated detectors with bandwidths above 67 GHz capable of handling baudrates of 100 Gbaud and more [7, 8]. We also demonstrated for the first time that a silicon-based EAM on a commercial platform could be more than suitable for this task with high bandwidths and good responsivity [9, 10].

Things becomes less obvious when looking at the typical bottlenecks in the optical modules today: the TIA, the CDR, and the modulator. Although the latest state-of-the-art TIAs have demonstrated 90 Gbaud capabilities at decent power levels for NRZ [11], 100 Gbaud PAM-4 will prove a challenging hurdle to take in the coming years.

Most 50 Gbaud optical modules today use either EMLs or SiPh MZMs as modulator. However, it seems unlikely that these MZMs will be able to double in bandwidth from approximately 35 GHz to 70 GHz by optimizing the current SiPh devices without significant help from analog or digital signal processing which will drastically affect the power consumption and cost, or without greatly reducing the optical power budget.

The topology introduced in Chapter 4 could provide a interesting alternative for these modulators, combining a compact, low-cost and low-power alternative. In [12], for example, we recently demonstrated 100 Gb/s PAM-4 transmission with this topology in combination with a dedicated 55 nm BiCMOS driver IC [13] at record-low powers of 1.5 pJ/bit (excluding the laser).

Even for the shortest data center links, disruptive changes might lay ahead in the not so distant future. As the data rates go up and more and more SiP based devices find their way to the data center market, single-mode optics will start to nibble at the dominance of the short (<50-100 m) MMF VCSEL-based links. These reaches will likely shrink even further as the increasing baudrates lead to much higher modal and chromatic distortion.

For now, SiP transceivers are more costly due to the relatively low volumes with respect to VCSELs-based transceivers and the need for a much more accurate (and thus more expensive) fiber alignment. Nevertheless, SiP transceivers could drop significantly in price with volume production, which is much less the case for competing platforms. Traditionally, VCSELs have proven unbeatable in terms of cost and have dominated the intra-rack interconnections. So it will be interesting to see if the VCSEL bandwidths can be engineered to sustain data rates at 53 Gbaud or higher without losing their cost benefit.

But perhaps the biggest question of all: will coherent optics get a foot on the ground inside the data center for the next-generation DCIs? Currently many module vendors are facing this issue: do we want the be early adopters and have the shortest time-to-market, but risk backing a solution that will not align with the future standard and thus possible fall out of favor? Or do we wait on the standardization efforts and risk losing a large part of the market due to a late entry?

In hindsight, the already difficult transition from the familiar and robust NRZ to PAM-4, will probably feel like a walk in the park in comparison with the transition from IMDD to coherent optics. Therefore, data center operators will most likely be hesitant to adopt coherent transceivers until they really need to. But if this work has taught me anything, it is this: it is best to tackle big tasks sooner rather than later.

References

[1] M. Theurer, H. Zhang, Y. Wang, W. Chen, L. Chen, Zeng, *et al.*, "2 × 56 gb/s from a double side electroabsorption modulated dfb laser and application in novel optical pam4 generation", *Journal of Lightwave Technology*, vol. 35, no. 4, pp. 706–710, 2017.

[2] A. Abbasi, L. Abdollahi Shiramin, B. Moeneclaey, J. Verbist, X. Yin, J. Bauwelinck, *et al.*, "Iii–v-on-silicon c-band high-speed electro-absorption-modulated dfb laser", *Journal of Lightwave Technology*, vol. 36, no. 2, pp. 252–257, Jan. 2018.

[3] A. Abbasi, J. Verbist, L. A. Shiramin, M. Verplaetse, T. de Keulenaer, R. Vaernewyck, *et al.*, "100-Gb/s Electro-Absorptive Duobinary Modulation of an InP-on-Si DFB Laser", *IEEE Photonics Technology Letters*, vol. 30, pp. 1095–1098, 2018.

[4] G. Roelkens, D. Vermeulen, S. K. Selvaraja, R. Halir, W. Bogaerts, and D. Thourhout, "Grating-based optical fiber interfaces for silicon-on-insulator photonic integrated circuits", *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 17, pp. 571–580, Jul. 2011.

[5] Albis Optoelectronics, *PD40H1: 56 Gbaud InGaAs Photodiode with Enhanced Responsivity*. [Online]. Available: `http://www.albisopto.com/albis_product/pd40f1-56-gbaud-photodiode-with-enhanced-responsivity/`.

[6] H. Ramon, M. Verplaetse, M. Vanhoecke, H. Li, J. Bauwelinck, P. Ossieur, *et al.*, "A 100-GS/s Four-to-One Analog Time Interleaver in 55-nm SiGe BiCMOS", *IEEE Journal of Solid-State Circuits*, pp. 1–1, 2021.

[7] H. Chen, P. Verheyen, P. D. Heyn, G. Lepage, J. D. Coster, S. Balakrishnan, *et al.*, ''-1 V bias 67 GHz bandwidth Si-contacted germanium waveguide p-i-n photodetector for optical links at 56 Gbps and beyond", *Opt. Express*, vol. 24, no. 5, pp. 4622–4631, Mar.

2016. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-24-5-4622.

[8] A. Beling, T. C. Tzu, J. Gao, J. S. Morgan, K. Sun, N. Ye, *et al.*, "High-speed integrated photodiodes", in *2019 24th OptoElectronics and Communications Conference (OECC) and 2019 International Conference on Photonics in Switching and Computing (PSC)*, 2019, pp. 1–3.

[9] J. Verbist, M. Verplaetse, S. A. Srivinasan, P. De Heyn, T. De Keulenaer, R. Pierco, *et al.*, "First Real-Time 100-Gb/s NRZ-OOK Transmission over 2 km with a Silicon Photonic Electro-Absorption Modulator", in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2017, pp. 1–3.

[10] J. Verbist, M. Verplaetse, S. A. Srinivasan, J. Van Kerrebrouck, P. De Heyn, P. Absil, *et al.*, "Real-time 100 gb/s nrz and edb transmission with a gesi electroabsorption modulator for short-reach optical interconnects", *Journal of Lightwave Technology*, vol. 36, no. 1, pp. 90–96, Jan. 2018.

[11] J. Lambrecht, H. Ramon, B. Moeneclaey, J. Verbist, M. Verplaetse, M. Vanhoecke, *et al.*, "90-gb/s nrz optical receiver in silicon using a fully differential transimpedance amplifier", *Journal of Lightwave Technology*, vol. 37, no. 9, pp. 1964–1973, May 2019.

[12] J. Lambrecht, J. Verbist, H. Ramon, M. Vanhoecke, J. Bauwelinck, X. Yin, *et al.*, "Low-Power (1.5 pJ/b) Silicon Integrated 106 Gb/s PAM-4 Optical Transmitter", *Journal of Lightwave Technology*, vol. 38, no. 2, pp. 432–438, Jan. 2020.

[13] H. Ramon, J. Lambrecht, J. Verbist, M. Vanhoecke, S. A. Srinivasan, P. De Heyn, *et al.*, "70 Gb/s Low-Power DC-Coupled NRZ Differential Electro-Absorption Modulator Driver in 55 nm SiGe BiCMOS", *Journal of Lightwave Technology*, vol. 37, no. 5, pp. 1504–1514, Mar. 2019.

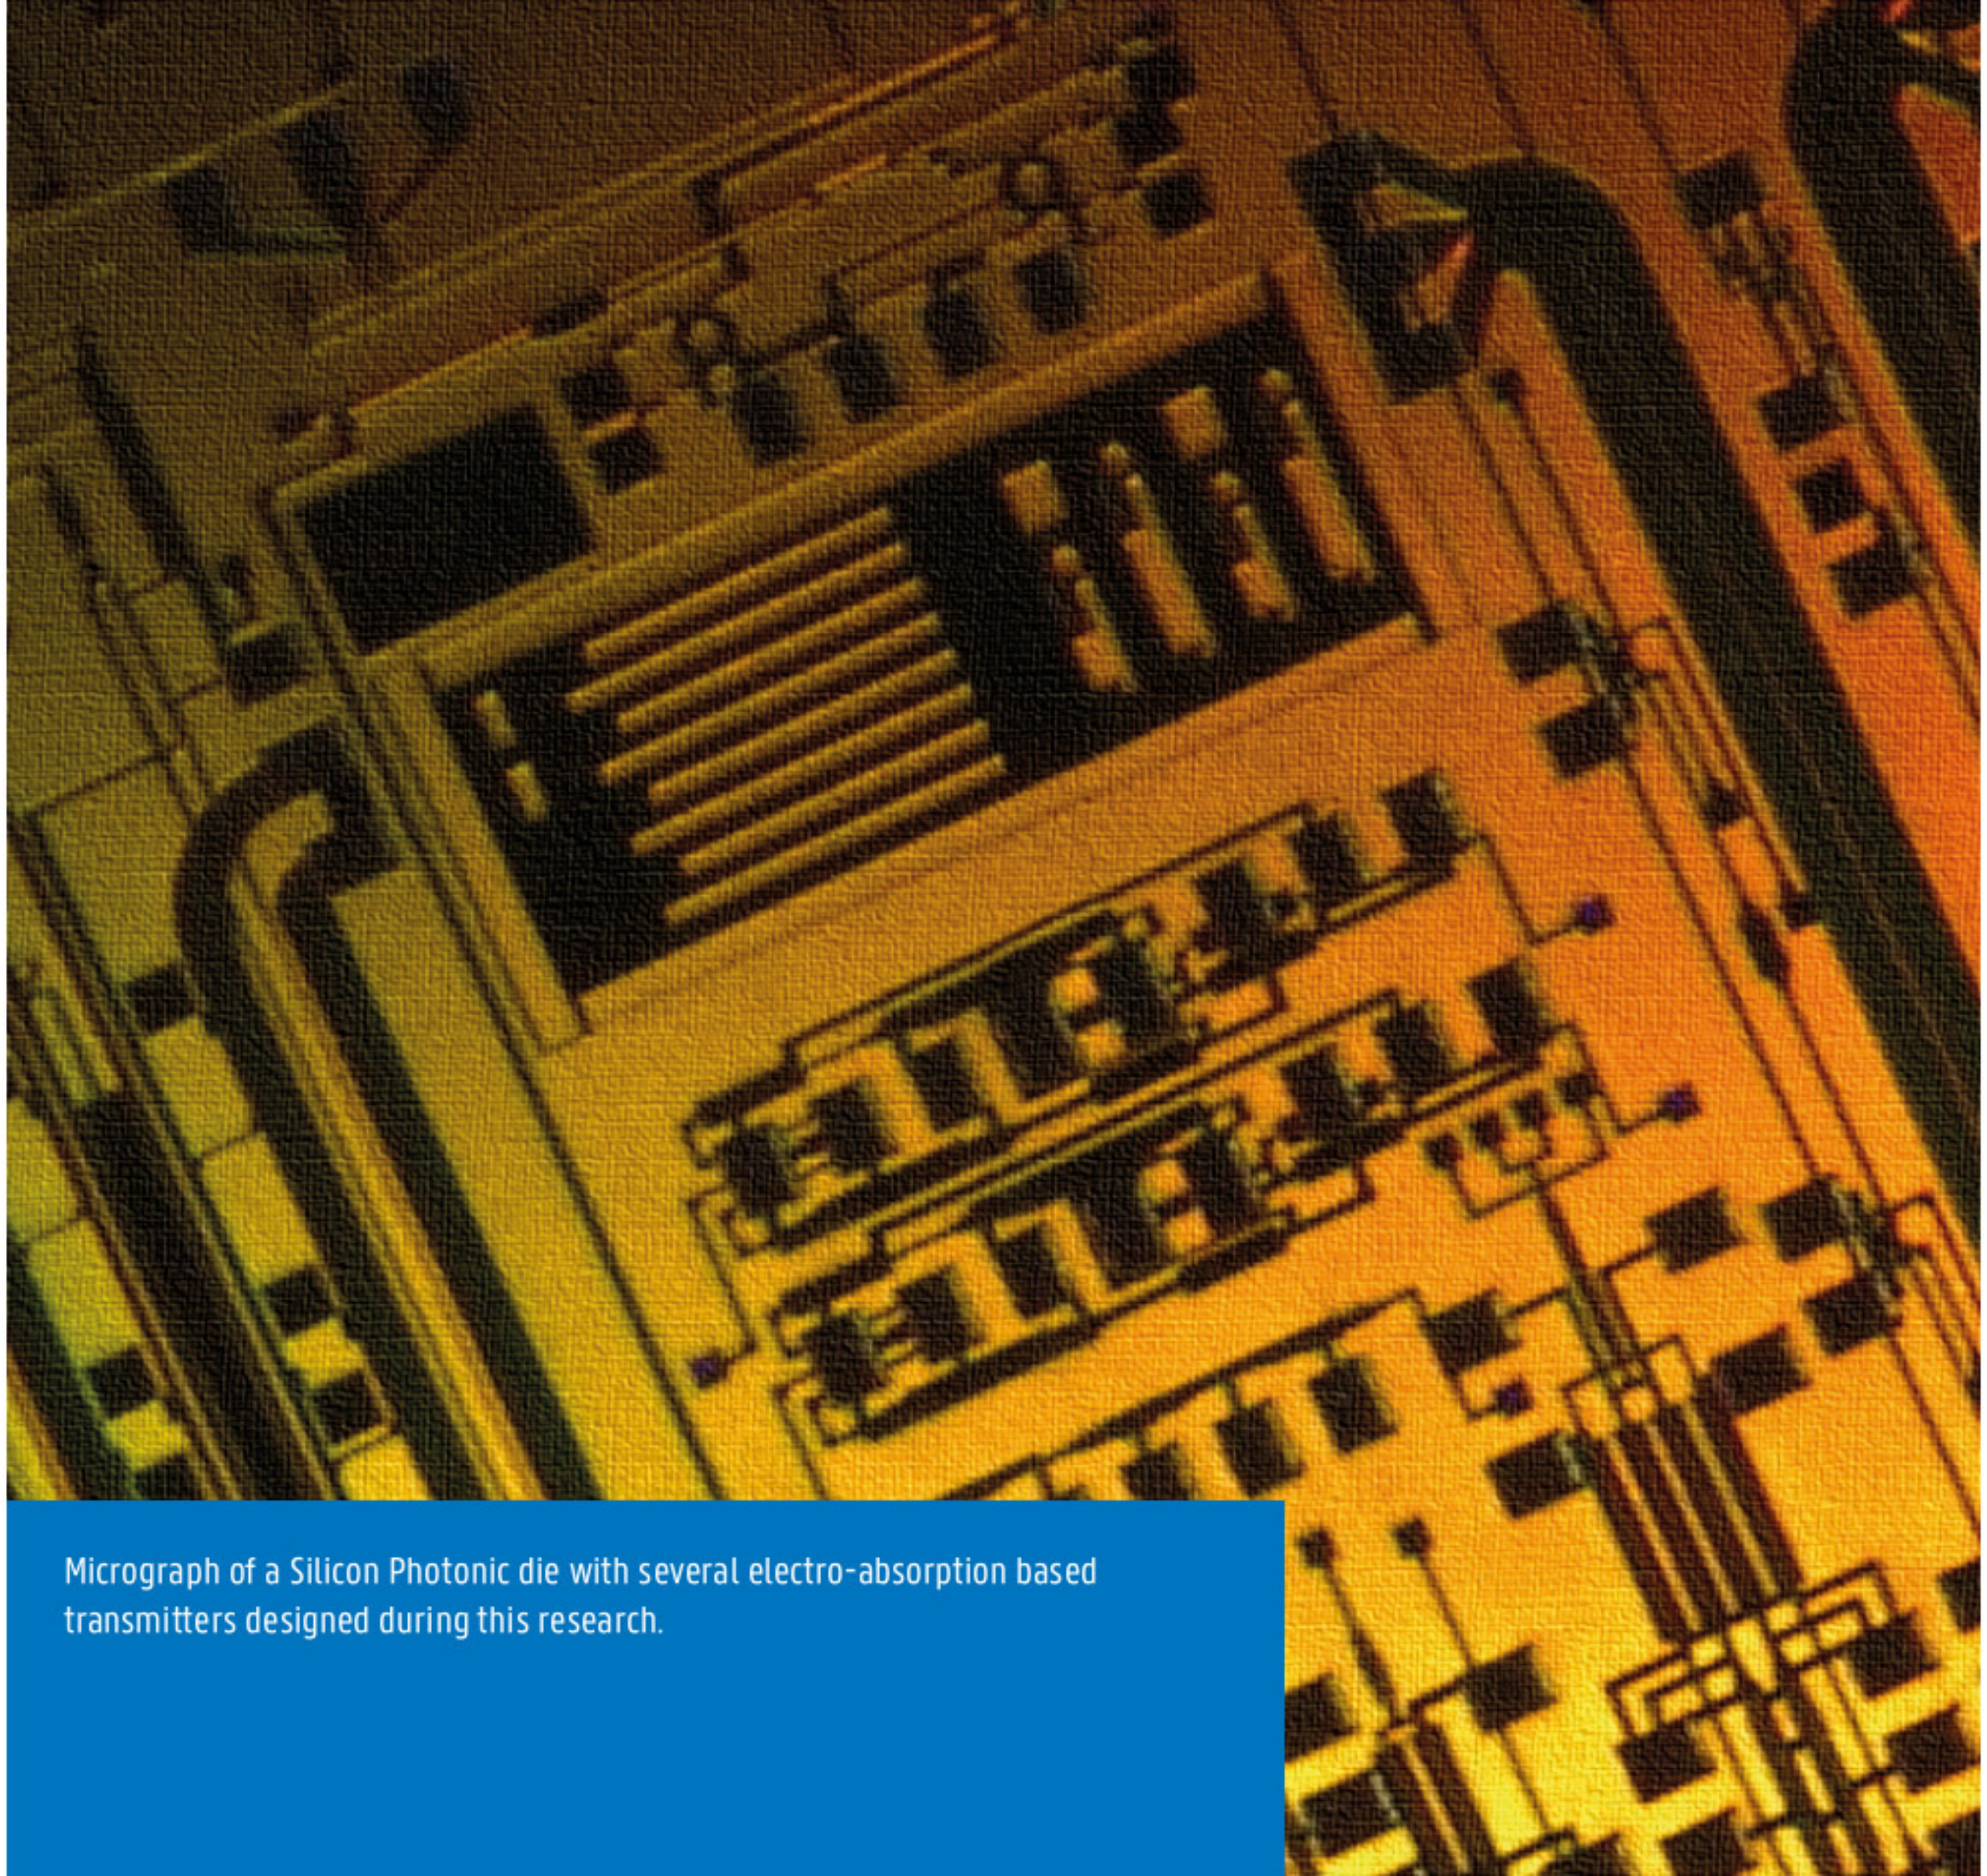

Micrograph of a Silicon Photonic die with several electro-absorption based transmitters designed during this research.