Een klok-en-data-herstelcircuit en MAC-voorverwerkingseenheid
aan meerdere snelheden met laag vermogen
voor 40 Gbit/s gecascadeerde bit-verwevende passieve optische netwerken

A Low Power, Multi-Rate Clock-and-Data Recovery Circuit and MAC Preprocessor
for 40 Gbit/s Cascaded Bit-Interleaving Passive Optical Networks

Arno Vyncke

UNIVERSITEIT
GENT

A Low Power, Multi-Rate Clock-and-Data Recovery Circuit and MAC Preprocessor for 40 Gbit/s Cascaded Bit-Interleaving Passive Optical Networks

Arno Vyncke

Publicly defended on November 16, 2016

Members of the examination board

prof. dr. ir. R. Van de Walle (Chairman)

INTEC, Ghent University, Belgium

prof. dr. ir. J. Bauwelinck (Secretary)

INTEC, Ghent University, Belgium

prof. dr. ir. X. Yin (Supervisor)

INTEC, Ghent University, Belgium

prof. dr. ir. G. Torfs (Supervisor)

INTEC, Ghent University, Belgium

prof. dr. ir. J. Doutreloigne

ELIS, Ghent University, Belgium

dr. ir. B. Lannoo

University of Antwerp, Belgium

prof. dr. L. Wosinska

KTH Royal Institute of Technology, Sweden

Dankwoord

Het is zover. Na wat een eeuwigheid leek, ligt het boek er. Toen ik in 2010 begon bij de Design onderzoeksgroep was het schrijven van een doctoraats-thesis iets wat in de verre toekomst lag. Maar de jaren glijden voorbij, elk jaar gevuld met interessante projecten, uitdagingen en nieuwe plannen. Zo-als men zegt, de tijd vliegt als je je amuseert. En het is door mij te amu-seren dat ik vandaag, 6 jaar later, met enige trots terugblik op het werk dat ik de afgelopen jaren verricht heb. Hoewel dit boek enkel het werk rond CBi-PON belicht, was dit niet mijn enige project. Door de verschillende projecten waarin ik betrokken raakte, kon ik mij verdiepen in de wereld van de digitale electronica – zowel FPGAs als ASICs. Ten slotte kreeg ik de kans om mijn bijdrage te leveren door de CABINET ASIC te ontwerpen in het kader van GreenTouch en het Europese DISCUS project, waardoor er ook analoog ASIC ontwerp aan te pas kwam.

Ik begon aan mijn avontuur met drie vrienden, die er voor zorgden dat de voorbije 6 jaar aangenaam doorgebracht konden worden. Legendarische Glee-nights, het doen verschijnen van BBQs op plaatsen waar ze nodig – maar niet gewenst – waren. We begonnen samen aan een klein avontuur in 2010, vandaag zijn we samen bezig aan een veel groter avontuur: de Bi-FAST spin-off. Renato, Timothy, Ramses, bedankt voor de voorbije jaren, het huidige jaar, en de komende jaren. Met jullie aan mijn zijde weet ik dat mij niet zal vervelen en dat we onze dromen kunnen waarmaken.

Door de omgeving waarin we werken veranderen de collega's vrij snel, maar gelukkig zijn er een aantal vaste waarden. Johan en Scott, bedankt voor jullie steun in ons verhaal, de leiding van het labo en het zorgen voor steeds weer nieuwe, uitdagende projecten. Guy, bedankt om samen de uit-daging die de CABINET was tot een goed einde brengen en waar nodig steeds erop te wijzen dat alles eigenlijk al opgelost is, op die eenvoudige implementatiestap na. Jean, jij was er bij vanaf mijn eerste contact met het

labo, tijdens ons fantastische VOP – waarvan ik nog steeds vind dat je je wil had moeten doordrijven – en bent doorheen de jaren toch ook net iets meer geworden dan zomaar een collega. Mike, bedankt om administratief alles in goede banen te leiden, ik ben blij dat ik het niet zelf moest doen. Ik wil ook onze vakgroepvoorzitter prof. De Zutter te bedanken voor de aangeboden faciliteiten en prof. Vandewege en prof. Qiu om het labo te brengen tot waar het vandaag staat.

Ik dien ook Christophe te bedanken, die mij met zijn engelengeduld in de wereld van de bit-interleaving PON introduceerde, en steeds klaarstond om zijn kennis en ervaring met mij te delen, zowel in het domein van de elektronica als daarbuiten. Jasmien, Jochen, Elber, Lucien en Elena, dank jullie voor de leuke jaren.

En dan zijn er nog de jonge, onbezonnen collega's, Bart, Koen, Wouter, Michaël, Marijn, Joris, Haolin, Gertjan, Manolo, Hubert, Hannes, Laurens, Michiel en José. Het is een plezier met jullie samen te werken en ik ben er van overtuigd dat jullie in de komende jaren voor nog veel mooie resultaten zullen zorgen.

Maar er is meer in het leven dan werk. Het zijn mijn vrienden die er voor zorgen dat ik de nodige rust vind in mijn hoofd wanneer nodig. Zij zijn het die vol interesse vragen waar ik mee bezig ben, er meestal weinig van snappen en dat ook helemaal niet erg vinden. Aan jullie allemaal: bedankt om een deel uit te maken van mijn leven.

Uiteraard wil ik ook mijn ouders, mama en papa, bedanken, die mij steeds alle kansen gegeven hebben om te doen waar ik interesse had en mij de brede, positieve kijk op het leven gegeven hebben waar ik dagelijks van geniet. Toendertijd misschien niet altijd met evenveel goesting, maar als ik er nu op terugblik ben ik blij al die ervaringen te mogen meegemaakt hebben. Ook mijn zus Elke verdient een woordje van bedanking, als grote zus heeft ze altijd goed voor mij gezorgd en dat probeert ze nog steeds wanneer nodig, met succes. Ook mijn schoonouders, Dirk en Hilde, verdienen hier de nodige erkenning. Jullie hebben mij de voorbije jaren opgenomen als jullie eigen zoon, en dat voelt goed. Er wordt beweerd dat schoonouders niet altijd een cadeau zijn, dus ik denk dat ik op dat vlak de hoofdvogel afgeschoten heb.

Elise, die de voorbije jaren alles van zeer dichtbij mocht meemaken. De lange werkuren door de passie waarmee ik met elektronica bezig was. De nog langere uren dat ik niet thuis was, omdat ik toch vond dat ik daarnaast ook een robot moest bouwen. Of aan een 24-uurs-programmeerwedstrijd moest deelnemen. Of omdat ik nu toch echt een eigen bedrijf moest op touw proberen zetten. Ik ben enorm blij dat we elkaar gevonden hebben, en kan me geen betere levensgezel inbeelden. Bedankt om me mijn dromen te laten volgen, en me daar steeds in te steunen. Bedankt voor je luisterend oor, al begrijpt het misschien niet altijd wat ik zeg. Bedankt om jouw dag met mij te delen.

Gent, november 2016

Arno Vyncke

"Sometimes I think the surest sign that intelligent life exists elsewhere in the universe is that none of it has tried to contact us."
Calvin (Bill Watterson)

Table of Contents

List of Figures

List of Tables

Acronyms

10GE-PON 10 Gbit/s Ethernet PON.

ADSL Asymmetric Digital Subscriber Line.

APON Asynchronous Transfer Mode (ATM) PON.

ASIC Application-Specific Integrated Circuit.

BAU Business-As-Usual.

BiPON Bit-interleaving PON.

BPON Broadband PON.

BWMAP Bandwidth Map.

CABINET CAscaded Bit-Interleaving eNd tErmination/repeaTer.

CapEx Capital Expenditure.

CBi-PON Cascaded Bit-interleaving PON.

CDR Clock-and-Data Recovery.

CML Current-Mode Logic.

CMOS Complementary Metal-Oxide-Semiconductor.

CO Central Office.

CP Charge Pump.

DBA Dynamic Bandwidth Allocation.

DLL Delay Locked Loop.

DNS Domain Name Service.

DSBWMAP Downstream Bandwidth Map.

DSL Digital Subscriber Line.

EFM Ethernet in the First Mile.

EPON Ethernet PON.

FLL Frequency Locked Loop.

FP7 Seventh Framework Programme.

FPGA Field-Programmable Gate Array.

FSAN Full Service Access Network.

FTTC Fiber-To-The-Curb.

FTTH Fiber-To-The-Home.

FTTP Fiber-To-The-Premises.

G-PON Gigabit-capable PON.

ICT Information and Communication Technology.

IEEE Institute of Electrical and Electronics Engineers.

IL Injection Locked.

IP Internet Protocol.

ISDN Integrated Services Digital Network.

ISP Internet Service Provider.

ITU International Communication Unit.

IXP Internet eXchange Point.

LA Limiting Amplifier.

LAN Local Area Network.

LD Laser Driver.

LR-PON Long-Reach PON.

MAC Medium Access Control.

NG-PON2 Next-Generation PON 2.

NNI Network-to-network interface.

NRZ Non-Return to Zero.

O-E-O Optical-Electrical-Optical.

ODN Optical Distribution Network.

OFDM Orthogonal Frequency Division Multiplexing.

OLT Optical Line Terminal.

P2MP Point-to-MultiPoint.

PC Personal Computer.

PD Phase Detector.

PI Phase Interpolator.

PLL Phase Locked Loop.

PN Phase Noise.

PON Passive Optical Network.

PWM Pulse-Width Modulation.

RN Remote Node.

RNID Repeater/End-ONT Identifier.

ROSA Receiver Optical Sub-Assembly.

SA-FF Sense Amplifier based Flip-Flop.

TDM Time Division Multiplexing.

TDMA Time Division Multiple Access.

TIA Transimpedance Amplifier.

TOSA Transmitter Optical Sub-Assembly.

TWDM Time and Wavelength Division Multiplexing.

UI Unit Interval.

USBWMAP Upstream Bandwidth Map.

VCO Voltage Controlled Oscillator.

VDSL2 Very-high-bit-rate Digital Subscriber Line 2.

WDM Wavelength Division Multiplexing.

WWW World Wide Web.

XG-PON 10-Gigabit-capable PON.

Nederlandse samenvatting
— Summary in Dutch —

Gedurende de laatste paar jaar is het dataverkeer exponentieel gestegen, en er wordt voorspeld dat het einde van deze groei nog niet in zicht is. Deze groei wijst op een steeds sterker verbonden wereld, wat vele voordelen met zich meebrengt: gemakkelijke toegang tot kennis, toename in productiviteit door efficiënte communicatie en het eenvoudig delen van rekenkracht. Deze voordelen hebben echter een prijs. De laatste jaren is er een sterke bewustwording opgekomen met betrekking tot de negatieve impact op het milieu die communicatienetwerken hebben als direct gevolg van hun enorme vermogenverbruik.

Onderzoek heeft aangetoond dat het vermogenverbruik van communicatienetwerken een significant en groeiend deel van het totale globale vermogenverbruik inneemt. Rekening houdend met de verbluffende groei die het dataverkeer nog steeds ondergaat, is het duidelijk dat de volgende-generatie netwerken een sterk verbeterde vermogensefficiëntie zullen moeten hebben indien we deze groei willen volhouden. Deze waarneming heeft geleid tot het oprichten van het GreenTouch consortium in 2010, die als missie had aan te tonen dat de energie-efficiëntie van communicatienetwerken verbeterd zou kunnen worden met een factor $1000\times$ tegen 2020, vergeleken met het door GreenTouch gedefinieerd *baseline netwerk* dat opgebouwd werd met de meest energie-efficiënte apparatuur beschikbaar in 2010.

Dit proefschrift stelt het Gecascadeerde Bit-verwevende Passief Optisch Netwerk (CBi-PON) voor, welke een van de vast-toegangsnetwerk technologieën ontwikkeld binnen GreenTouch is. De CABINET ASIC wordt geïntroduceerd als een instantiatie van een generiek CBi apparaat dat gebruikt kan worden om een CBi-PON uit te rollen. Vermogenmetingen op de CABINET laten ons toe de vermogenverbruiksreductie te schatten van een volgende-generatie netwerk dat een CBi-PON gebruikt in plaats van een meer traditioneel Passief Optisch Netwerk. Het onderzoek voor dit proef-

xxiv

schrift werd uitgevoerd in het kader van het GreenTouch CBI project en het
EU FP7 DISCUS project, beide projecten die technologieën ontwikkelen
die moeten leiden tot duurzame communicatienetwerken.

In Hoofdstuk 1 wordt de lezer bekend gemaakt met de oorsprong van het
Internet, om zo de nodige context te verschaffen waarbinnen het werk voor
dit proefschrift werd uitgevoerd. Het hoofdstuk vervolgt met het onthullen
van de problemen met de huidige communicatienetwerken, waarbij zowel
de ecologische als de economische impact besproken wordt.

Vervolgens wordt de huidige *core-metro-access* architectuur voorgesteld en
worden typische getallen gegeven om aan te tonen waarom het vermogen-
verbruik van de *access* en *metro* delen het leeuwendeel van het totale ver-
mogenverbruik voor zich nemen, te wijten aan het enorme aantal apparaten
in het netwerk. Het access netwerk wordt besproken, waarbij de beper-
kingen van koper, de verschuiving naar optische toegangsnetwerken en het
concept van passieve optische netwerken worden behandeld. De metro-
access convergentie door de introductie van *Long-Reach* PONs wordt aan-
gehaald als de meest recente ontwikkeling. De uitdagingen voor volgende-
generatie netwerken worden besproken, waaruit de twee belangrijkste uit-
dagingen voortkomen: (1) hoge bandbreedtes per gebruiker en (2) flexibele
bandbreedte-allocatie, beiden gecombineerd met laag vermogenverbruik.

De Bit-verwevende PON (BiPON), die de basis vormt waarop de Gecasca-
deerde Bit-verwevende PON gebouwd is, wordt kort voorgesteld in Hoofd-
stuk 2. Het hoofdstuk start met een overzicht van de verschillende PON
protocollen, om zo de lezer te voorzien van het referentiekader waarbinnen
BiPON ontwikkeld werd. Daaropvolgend wordt de belangrijke paradigma-
verschuiving van pakket-gebaseerde *Time Division Multiplexing* (TDM)
naar bit-gebaseerde TDM uitgelegd en wordt de ingebouwde dynamische
bandbreedte-allocatie toegelicht. Het hoofdstuk eindigt met het presenteren
van de enorme vermogensreductie van $35\times$ tot $180\times$ die te danken is aan
het gebruik van BiPON.

Hoofdstuk 3 bespreekt de metro-access convergentie en Long-Reach PONs
als een oplossing voor volgende-generatie netwerken. Een van de moge-
lijke problemen is de impact op de prijs en het vermogenverbruik van de
Optische Netwerk Eenheid (ONU) bij de eindgebruiker. Deze kan aange-
pakt worden door BiPON tot het volgende niveau te verheffen door meer-
dere, gecascadeerde niveaus van BiPON te implementeren: de Gecasca-
deerde Bit-verwevende PON (CBi-PON). Het hoofdstuk vervolgt met de
introductie van de CBi netwerktopologie en de verschillende CBi appara-

ten nodig in een CBi-PON: CBi Interleaver, CBi Repeater en CBi End-ONT. Details worden gegeven rond de compositie van het CBi Frame en een 3-niveau instantiatie van een CBi-PON wordt gepresenteerd. Deze instantiatie gebruikt een generiek CBi apparaat welke zowel de rol van CBi Repeater, als die van CBi End-ONT kan opnemen: de CABINET.

Het ontwerp en de implementatie van de CABINET ASIC wordt besproken in Hoofdstuk 4. De systeemarchitectuur wordt voorgesteld en deze bestaat uit drie grote delen: (1) de analoge front-end, (2) de Medium Access Control (MAC) voorverwerkingseenheid en (3) de analoge back-end.

De analoge front-end bestaat uit een klok-en-data herstelcircuit (CDR) aan meerdere snelheden. De topologiekeuze van de CDR wordt verklaard en de keuzes worden toegelicht die toelaten om de CDR te ontwerpen in de geest van BiPON: de snelheid zo vroeg mogelijk in de verwerkingsketen reduceren. Bijgevolg is de 40 Gbit/s CDR een *sub-sampling* CDR, wat wil zeggen dat deze niet alle databits herstelt. De kritische bouwblokken worden behandeld: de fasedetector, bemonsteraar, spanningsgestuurde oscillatoren (40 GHz, 10 GHz en 2.5 GHz), *charge pump* en lusfilter. De dimensionering wordt besproken en de simulatieresultaten tonen aan dat alle CDRs *locken* zoals verwacht.

Volgend op de analoge front-end wordt de MAC-voorverwerkingseenheid besproken. De MAC-voorverwerkingseenheid is een puur digitale blok, welke de herstelde databits van de analoge front-end neemt en verwerkt volgens het CBi-PON protocol om zo de correcte data te bepalen zodat deze doorgestuurd kan worden (CBi Repeater) of naar een FPGA gezonden kan worden voor verdere verwerking (CBi End-ONT). In het geval van de CBi Repeater serialiseert de analoge back-end de uitgang van de MAC-voorverwerkingseenheid tot een enkele datastroom die doorheen het lagere CBi niveau gestuurd kan worden. De details van de implementatie van deze analoge back-end worden gegeven op het einde van Hoofdstuk 4. Het hoofdstuk sluit af met een visualisatie van de CABINET ASIC layout.

In het eerste deel van Hoofdstuk 5 worden de meetresultaten van de CABINET ASIC voorgesteld. De spanningsgestuurde oscillatoren (VCO's) zijn gemeten en er is aangetoond dat ze werken zoals verwacht - al is hun gevoeligheid aan de voedingsspanning hoger dan geanticipeerd. Daarenboven werden er problemen waargenomen met de *serializer*, vermoedelijk veroorzaakt door problemen met de klokdistributie. De *Frequency Locked Loop* (FLL) werkt zoals verwacht voor alle snelheden. Hoewel de bemonsteraars de mogelijkheid tonen om een enkele bit te bemonsteren aan

28 Gbit/s, heeft de bemonsterklok die afgeleid is van de VCO's te veel jitter om een goede bemonstering te bekomen. Bijgevolg geven de metingen op de CDR-werking van de CABINET geen bevredigende resultaten. Gebruikmakend van dataherstel op basis van de FLL waren we wel in staat 2.5 Gbit/s datastromen correct te herstellen en de werking van de MAC-voorverwerkingseenheid werd geverifieerd op deze datastromen. Uit de vermogensmetingen van deze correct werkende modus werden vermogenverbruiksschattingen voor alle modi geëxtrapoleerd.

In het tweede deel van het hoofdstuk worden de geëxtrapoleerde schattingen van het vermogenverbruik aangewend om de reductie in vermogenverbruik te berekenen die de toepassing van een CBi-PON kan betekenen. De referentie waartegen het vermogenverbruik vergeleken wordt is het *2010 baseline netwerk* zoals gedefinieerd door het GreenTouch consortium. Het vermogenverbruik van een ONU daalt van 1.5 W tot 0.4 W, of een vermogenverbruik-reductiefactor van 3.7$\times$. De Remote Node gebruikt in CBi-PON verbruikt slechts 21.17 mW/ONU en vervangt de sterk verbruikende aggregatieswitch die 121 mW/ONU verbruikt, wat zich vertaalt in een vermogenverbruik-reductiefactor van 5.7$\times$.

Dit proefschrift wordt afgesloten met Hoofdstuk 6, waarin het werk wordt samengevat, een vooruitzicht wordt gegeven op volgende-generatie netwerken en een toelichting wordt gegeven over het werk dat verder nog kan gebeuren in de context van CBi-PON en volgende-generatie netwerken in het algemeen.

English summary

During the last couple of years, data traffic has been rising exponentially and it is predicted that this growth is not going to end anytime soon. This growth indicates a more connected world which brings numerous benefits: easier access to knowledge, productivity increases due to efficient communication and facilitated sharing of computing power. However, these benefits come with a price. Over the past few years a stronger awareness has risen with respect to the negative environmental impact communication networks have as a direct result of their massive power consumption.

Research has shown that the power consumption of communication networks is taking up a significant and growing share of the total global power consumption. Taking into account the staggering growth data traffic is still undergoing, it is clear next generation networks are going to require a much improved power efficiency if we want to sustain this growth. This observation led to the foundation of the GreenTouch consortium in 2010, which made it its mission to show that the energy efficiency of communication networks could be improved by a factor of $1000\times$ by 2020, compared to the GreenTouch-defined baseline network which was built using the most energy efficient equipment available in 2010.

This dissertation presents the Cascaded Bit-interleaving PON (CBi-PON), which is one of the Fixed Access Network technologies developed within GreenTouch. The CABINET ASIC is introduced as an instantiation of a generic CBi Device that can be used to deploy a CBi-PON. Power measurements on the CABINET allow us to estimate the power consumption reduction that can be expected from a next generation network that uses a CBi-PON instead of the more traditional Passive Optical Networks. The research conducted for this dissertation was performed as part of the GreenTouch CBI project and the EU FP7 DISCUS project, both projects that aim to develop technologies resulting in sustainable communication networks.

In Chapter 1, the reader is introduced to the origins of the Internet to provide the required context in which the work for this dissertation was performed.

The chapter continues to reveal the issues with current communication networks, highlighting both the ecological as the economical impact.

Subsequently, the current *core-metro-access* architecture is presented and typical numbers are given to show why the power consumption of *access* and *metro* tiers constitute the lion's share of the total power consumption due to the vast amount of devices in the network. The access network is discussed, covering the copper limitations, the move to all-optical access networks and the concept of passive optical networks. Finally, the metro-access convergence with the introduction of Long-Reach PONs is presented as the latest development. The challenges of next-generation networks are discussed, arriving at two main challenges: (1) high bandwidths per user and (2) flexible bandwidth allocation, both combined with lower power consumption.

The Bit-interleaving PON (BiPON), which forms the basis on which the Cascaded Bit-interleaving PON is built, is briefly introduced in Chapter 2. The chapter starts off with an overview of the different PON protocols, in order to provide the reader with the necessary reference frame in which BiPON was developed. Next, the important paradigm-shift from packet-based Time Division Multiplexing (TDM) to bit-based TDM is explained and the built-in dynamic bandwidth allocation is highlighted. The chapter ends by presenting the enormous power savings of $35\times$ up to $180\times$ associated to the use of BiPON.

Chapter 3 discusses the metro-access convergence and Long-Reach PONs as a solution for next-generation networks. One of the possible issues is the impact on the cost and the power consumption of the Optical Network Unit (ONU). This can be mitigated by taking BiPON to the next level by implementing multiple, cascaded levels of BiPON: the Cascaded Bit-interleaving PON (CBi-PON). The chapter continues to introduce the CBi Network Topology and the different CBi Devices required in a CBi-PON: CBi Interleaver, CBi Repeater and CBi End-ONT. Details are given on the CBi Frame composition and finally a 3-level instantiation of a CBi-PON is presented. This instantiation uses a generic CBi Device which can serve both as CBi Repeater and CBi End-ONT: the CABINET.

The CABINET ASIC design and implementation is discussed in Chapter 4. The system architecture is presented and three main parts are identified: (1) the Analog Front-End, (2) the Medium Access Control (MAC) preprocessor and (3) the Analog Back-End.

The Analog Front-End consists of a multi-rate Clock-and-Data Recovery (CDR) circuit. The choice of CDR topology is clarified, as well as the choices made to design the CDR in the spirit of the Bit-interleaving PON: reducing the rate as early in the chain as possible. As a result, the 40 Gbit/s CDR is a sub-sampling CDR, which means it is not recovering all data bits. All critical building blocks are covered: the Phase Detector, Sampler, Voltage-Controlled Oscillators (40 GHz, 10 GHz and 2.5 GHz), Charge Pump and Loop Filter. The sizing is discussed and simulation results show that all CDRs lock as expected.

Following the Analog Front-End, the MAC preprocessor is discussed. The MAC preprocessor is a purely digital block, which takes in the recovered data bits from the Analog Front-End and processes them according to the CBi-PON protocol to determine the correct data to either forward (CBi Repeater) or send to an FPGA for further processing (CBi End-ONT). In the case of a CBi Repeater, the Analog Back-End serializes the MAC preprocessor output to a single data stream that can be sent through the lower CBi Level. The details of the implementation of this Analog Back-End are given at the end of Chapter 4. The chapter is concluded with a visualization of the CABINET ASIC layout.

In the first part of Chapter 5, the measurement results of the CABINET ASIC are presented. The VCOs are measured and shown to operate as expected - although their power supply sensitivity is higher than anticipated. Furthermore, issues were detected with the serializer, most likely caused by the clock distribution. However, the Frequency Locked Loop (FLL) is verified to operate as expected for all rates. While the samplers show the ability to sample a single bit at 28 Gbit/s, the sampling clock which is derived from the VCOs contains too much jitter to obtain any good data sampling. As a result, the measurements on the CDR operation of the CABINET could not produce satisfying results. Using FLL-based data recovery, we were able to recover 2.5 Gbit/s data streams correctly and the MAC preprocessor was verified to operate as expected for these data streams. From the power measurements on this correctly operating mode, power measurements for all modes were extrapolated.

In the second part of the chapter, the extrapolated power estimates are used to calculate the power consumption reduction which the use of a CBi-PON could bring. The reference to which the power consumption is compared is the 2010 baseline network as defined by the GreenTouch consortium. Comparing the power consumption of an ONU, which drops from 1.5 W

to 0.4 W, a power consumption reduction factor of $3.7\times$ is shown. The Remote Node used in the CBi-PON consumes only 21.17 mW/ONU and replaces the power-hungry Aggregation Switch that consumes 121 mW/ONU, which translates to a power consumption reduction factor of $5.7\times$.

The dissertation is concluded with Chapter 6, which summarizes the work performed, gives an outlook on next-generation networks and highlights future work that could be done in the context of CBi-PON and next-generation networks in general.

List of publications

Publications in international journals

- **A. Vyncke**, G. Torfs, C. Van Praet, M. Verbeke, A. Duque, D. Su-vakovic, H. Chow, and X. Yin, *The 40 Gbps Cascaded Bit-interleaving PON [invited]*, Optical Fiber Technology

- M. Verbeke, P. Rombouts, **A. Vyncke**, and G. Torfs, *Influence of jitter on limit cycles in bang-bang clock and data recovery circuits*, IEEE Transactions on Circuits and Systems I - Regular Papers, vol. 62, no. 6, July 1, 2015, pp. 1463-1471

- T. De Keulenaer, G. Torfs, Y. Ban, R. Pierco, R. Vaernewyck, **A. Vyncke**, Z. Li, J. H. Sinsky, B. Kozicki, X. Yin, and J. Bauwelinck, *84 Gbit/s SiGe BiCMOS duobinary serial data link including seri-aliser/deserialiser (SERDES) and 5-tap FFE*, IET Electronic Letters, vol. 51, no. 4, February 19, 2015, pp. 343-345

- X. Z. Qiu, X. Yin, J. Verbrugghe, B. Moeneclaey, **A. Vyncke**, C. Van Praet, G. Torfs, J. Bauwelinck, and J. Vandewege, *Fast synchroniza-tion 3R burst-mode receivers for passive optical networks [invited tutorial]*, Journal of Lightwave Technology, vol. 32, no. 4, February 15, 2014, pp. 644-659

- C. Van Praet, G. Torfs, **A. Vyncke**, E. Matei, P. Cautereels, and J. Bauwelinck, *Fast H.264 intra prediction for network video process-ing*, IEICE Electronics Express, vol. 10, no. 12, June 2, 2013, pp. 1-6

Publications in international conferences

- J. Van Kerrebrouck, T. De Keulenaer, J. De Geest, R. Pierco, R. Vaernewyck, **A. Vyncke**, M. Fogg, M. Rengarajan, G. Torfs, and

J. Bauwelinck, *100 Gb/s serial transmission over Copper using duo-binary signaling*, Designcon, Santa Clara, California, USA, January 19-21, 2016

- **A. Vyncke**, G. Torfs, M. Verbeke, C. Van Praet, H. Chow, D. Suvakovic, A. Duque, and X. Yin, *A Low Power 40 Gbit/s Cascaded Extension to Bit-Interleaving Optical Networks Enabling Next-Generation Metro/Access Connectivity*, 20th Annual Symposium of the IEEE Photonics Benelux Chapter, Brussels, Belgium, November 26-27, 2015

- **A. Vyncke**, G. Torfs, M. Verbeke, C. Van Praet, H. Chow, D. Suvakovic, A. Duque, and X. Yin, *Voltage controlled oscillators for 40 Gbit/s cascaded bit-interleaving PON*, Advances in Wireless and Optical Communications (RTUWO 2015), 5-6 November 2015, Riga, Latvia, November 5-6, 2015, pp. (1-4)

- **A. Vyncke**, G. Torfs, M. Verbeke, and X. Yin, *An 8-phase 10 GHz Voltage Controlled Ring Oscillator for 40 Gbit/s BiPON Clock-and-Data Recovery*, 11th Conference on PhD Research in Microelectronics and Electronics (IEEE PRIME 2015), Glasgow, United Kingdom, June 29 - July 2, 2015, pp. (1-4)

- **A. Vyncke**, G. Torfs, M. Verbeke, C. Van Praet, H. Chow, D. Suvakovic, A. Duque, and X. Yin, *CBI-PON: a Low Power Solution offering Flexible Bandwidth Allocation for 40 Gbit/s Next generation Metro/Access Networks*, IEICE Information and Communication Technology Forum 2015 (ICTF 2015), Manchester, United Kingdom, June 3-5, 2015, pp. (1-4)

- X. Yin, H. Chow, **A. Vyncke**, D. Suvakovic, G. Torfs, A. Duque, D. van Veen, M. Verbeke, T. Ayhan, and P. Vetter, *CBI: a scalable energy-efficient protocol for metro/access networks [invited]*, 2014 IEEE Online Conference on Green Communications (OnlineGreenComm), November 12-14, 2014, pp. 126-131

- G. Torfs, X. Yin, **A. Vyncke**, M. Verbeke, J. Bauwelinck, *Solutions for a single carrier 40 Gbit/s downstream long-reach passive optical network [invited]*, 16th International Telecommunications Network Strategy and Planning Symposium (Networks 2014), vol. fri.20.2, Funchal, Madeira, Portugal, September 17-19, 2014, pp. 1-5

- X. Yin, X. Z. Qiu, G. Torfs, C. Van Praet, R. Vaernewyck, **A. Vyncke**, J. Verbrugghe, B. Moeneclaey, M. Ruffini, D. B. Payne, and J. Bauwelinck, *Performance evaluation of single carrier 40-Gbit/s downstream for long-reach passive optical network*, 18th International Conference on Optical Network Design and Modeling (ONDM 2014), Stockholm, Sweden, May 19-22, 2014, pp. 162-167

- J. Bauwelinck, R. Vaernewyck, J. Verbrugghe, W. Soenen, B. Moeneclaey, C. Van Praet, **A. Vyncke**, G. Torfs, X. Yin, X. Z. Qiu, J. Vandewege, N. Sotiropoulos, H. de Waardt, R. Cronin, G. Maxwell, T. Tekin, P. Bakopoulos, C. P. Lai, and P. D. Townsend, *High-speed electronics for short-link communication [invited]*, 39th European Conference and Exhibition on Optical Communication (ECOC 2013), vol. Mo.4.F.4, London, UK, September 22-26, 2013, pp. 164-166

- W. Soenen, R. Vaernewyck, **A. Vyncke**, and J. Bauwelinck, *Evaluation of a discrete 4-PAM optical link for future automotive networks*, Annual Symposium of the IPS Benelux Chapter, pp. 69-72, November, 29-30, 2012, Mons, Belgium

Publications in national conferences

- **A. Vyncke**, *Digital Assistant Processor Relieves Analog Designer Headaches*, 13th FEA PhD Symposium, December 5, 2012, Ghent, Belgium

Chapters in books

- **A. Vyncke**, G. Torfs, M. Verbeke, C. Van Praet, H. Chow, D. Suvakovic, A. Duque, and X. Yin, *Design and measurement of VCOs for 40 Gbit/s Cascaded Bit-Interleaving PON*, 1st International IEEE Conference on Advances in Wireless and Optical Communications 2015, Latvia, Riga, 5-6 November, 2015. Riga: RTU Press, 2015, pp.91-104. ISBN 978-9934-10-758-0

1

Introduction

Mankind has always displayed a strong desire to communicate. Languages were invented, writing was developed and for years these were the main means of communication. By the second half of the 19th century, scientific progress enabled an extraordinary pace at which new communication technologies became available: printing, photography, telegraphy, telephony, ... With the introduction of the Internet, an unprecedented convenience in communication was born. Providing a number of virtually unlimited direct connections to anywhere in the world, the Internet fulfills mankind's need for communication like no technology ever before.

Remarkably, even though the Internet has become such an important part of everyday life, there is a general ignorance regarding the infrastructure allowing the Internet to achieve the impressive world-wide connectivity that is provided. As this infrastructure forms the context in which the research in this dissertation has been conducted, this introductory chapter starts off with a brief history of how the Internet originated. Subsequently, the current form of the Internet's infrastructure is described and this chapter is concluded by discussing next-generation networks and the challenges these will be facing.

1.1   On the origins of Internet

Even though most of us use the Internet on a daily basis, few really know the underlying infrastructure that enables us to send each other messages, instantly see news from all around the world or stream movies at home in high definition quality. Despite its current ubiquitousness, the road to today's world-wide network was not an easy one. In what follows, a brief overview is given on how the Internet originated and how our communication networks are organized today.

1.1.1   From ARPANET to the Internet - a brief history

The birth of the Internet took place during the Cold War. At that moment in time there was only one network available in the United States: the AT&T telephone network. The design of the AT&T network used a centralized architecture, making it extremely vulnerable to malfunctions or attacks. Of course, the risk of losing all means of communication while under a constant threat of a nuclear attack was not a favorable position to be in.

Moreover, at that time, DARPA (Defense Advanced Research Projects Agency - the US Defense Research Department) had acquired multiple research computers on different geographical locations. Enabling communication between these computers was considered advantageous, since it would make the capabilities of a computer installed at one site accessible by a research group located at a different one. Furthermore, the possibility to share data would significantly improve collaboration between the different research groups.

To address these needs, DARPA developed the first ever packet-switching network, setting up a connection between four different nodes: the UCLA (University of California, Los Angeles), SRI (Stanford Research Institute), UCSB (University of California, Santa Barbara) and the University of Utah (Figure 1.1). This network, coined ARPANET, became a reality in September 1969 with the first successful transmission of data between the UCLA and SRI, and is considered the precursor of the Internet. However, it was still very different from the Internet as we know it today. As the number of sites connected to ARPANET quickly grew, the TCP/IP protocol was developed in 1983 to handle compatibility issues on the network. This allowed more computers to easily connect to ARPANET, leading to an even more rapid growth.

In 1986 the National Science Foundation (NSF) had created a couple of supercomputer centers at several universities, and the NSF introduced NSFNet
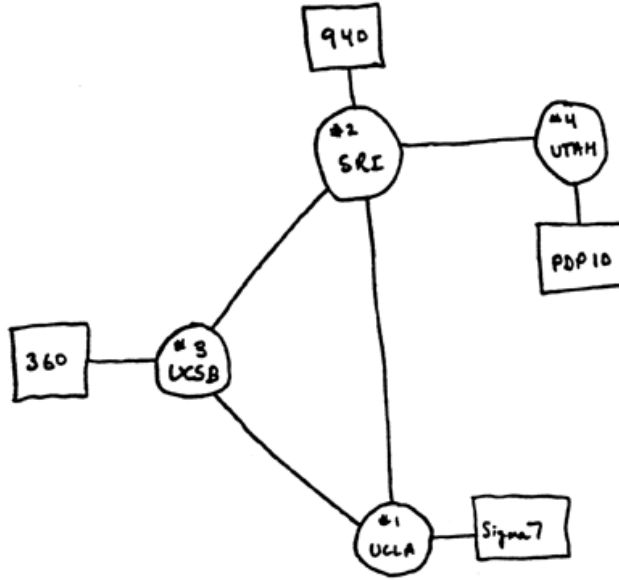
Figure 1.1: First ARPANET with 4 nodes [1]
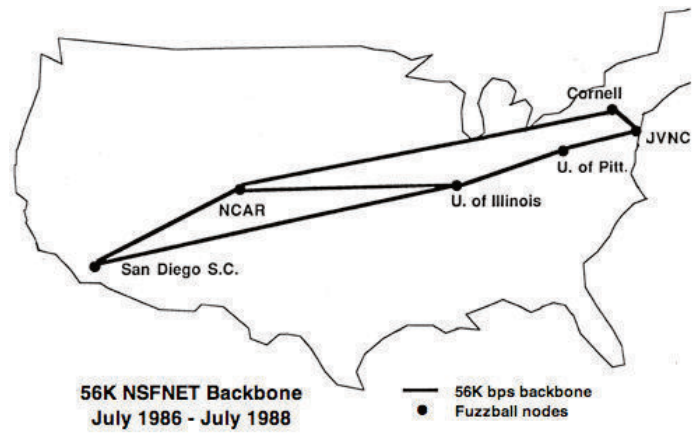
Figure 1.2: Original NSFNet with 56 kbit/s backbone [2]

as a way to connect the different supercomputer sites, as shown in Figure 1.2. In the same year, Berkeley released BSD (Berkeley Software Distribution) Unix, a multi-user operating system including a complete version of TCP/IP (Transmission Control Protocol/Internet Protocol).

With the release of BSD Unix, which was free, colleges and universities started to create local area networks. Since they used TCP/IP, they could connect these smaller networks to NSFNet, effectively creating an internetwork: a network consisting of networks, where NSFNet was functioning as the so-called *backbone* of the internetwork.
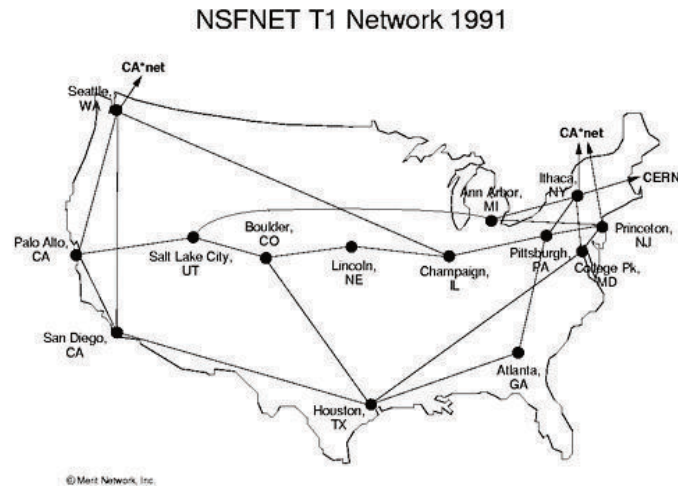
Figure 1.3: Intermediary NSFNet with T1 backbone (1.5 Mbit/s) [3]

During this time, the NSFNet backbone had a bandwidth of 56 kilobits per second. Its popularity grew tremendously, requiring NSF to upgrade the NSFNet backbone speed to 1.5 Mbit/s (T1) by 1988, which was only 2 years after its conception. The growing popularity of NSFNet can partly be explained by the introduction of the Domain Name Service (DNS) system in 1984. Until then, computers were addressed by their Internet Protocol (IP) address, which is difficult to remember (e.g. 74.125.228.194). The DNS system can be thought of as an automatic phonebook, allowing one to use an easy to remember name (e.g. google.com), which is then translated by the DNS system to the correct IP address. Eliminating the need to memorize complicated numbers definitely facilitated the adoption of NSFNet by the broader public.

Since ARPANET did not receive a bandwidth upgrade, it was no longer

NSFNET T3 Network 1992
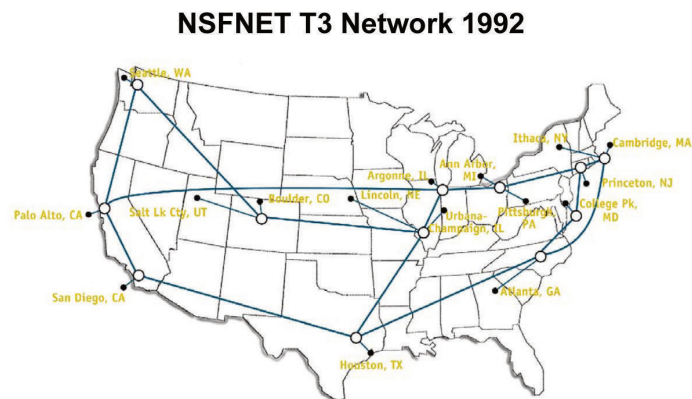
Figure 1.4: Final NSFNet with T3 backbone (45 Mbit/s) [4]

relevant, causing it to be disconnected from the Internet in 1990. NSFNet on the other hand had to be upgraded again in 1991 (Figure 1.3) up to 45 Mbit/s (T3) to support the ever-increasing data traffic. The NSFNet with the T3 backbone as it was in 1992 is visualized in Figure 1.4.

The introduction of several commercial networks, which were interconnected using the federal-funded, academic-oriented NSFNet caused a disruption in the philosophy regarding how the Internet should be organized. As a result, the NSFNet backbone service was shut down in 1995 and replaced by a multitude of commercial backbone service providers. Network Access Points provided traffic exchanges between the different commercial networks, a task that later would be performed by Internet eXchange Points (IXPs).

At this point in time, other countries had already connected their networks to the US network, making the Internet an international concept. The interconnection between the different backbone networks was provided by IXPs, which is still the case today.

Today, so many commercial Internet Service Providers exist, that they are categorized in 3 different tiers which are hierarchically interconnected, as shown in Figure 1.5. A Tier 1 Internet Service Provider (ISP) operates at the highest level, and only a few exist. A Tier 1 has an extensive network, covering a large area. Tier 1 ISPs typically have so called *peering agreements*: agreements that define how Tier 1 ISPs make use of each other's network to provide world-wide connectivity to their customers.

Compared to a Tier 1 ISP, a Tier 2 ISP has a more limited coverage span and
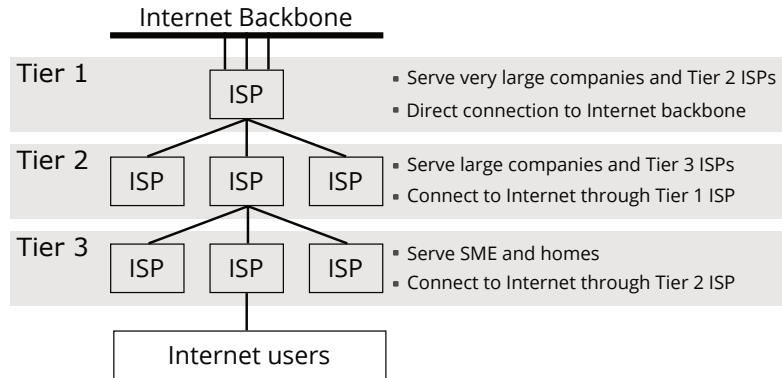
Figure 1.5: Internet Service Provider Tiers

needs to buy Tier 1 capacity to route the data traffic of its customers over the Tier 1 network that provides the desired global connectivity. Finally, at the lowest level are the Tier 3 ISPs. These typically operate at regional level, such as Proximus and Telenet in Belgium, and deliver Internet connectivity to end-users. Of course, these Tier 3 ISPs have contracts in place with either Tier 1 or Tier 2 ISPs to provide that service to their end-customers.

Historically, Internet traffic was routed over existing telephone copper cables. Those were once installed when deploying the telephone network and were readily available. Therefore it was cost-effective to re-use these. However, due to the rapid growth of the Internet, the data traffic increased exponentially. This urged ISPs to upgrade parts of the network to support higher bandwidths. For a few years new techniques allowed to keep using copper cable, but in the end optical fiber had to be installed to provide the needed bandwidths.

Because of the hierarchical nature of the networks, it was logical to upgrade the core network to optical fiber first, while the copper cable access network remained untouched. Throughout the years, more and more parts of the access network have been replaced by optical fiber, up to the point that we can talk about *the last mile*: only the part of the network that connects your house to the first aggregation switch is still in copper. That means a large share of the legacy copper network has already been replaced by optical fiber today. However, since optical fiber was installed merely as a drop-in replacement for copper cable, the network as a whole does not exploit all advantages of optical fiber. Therefore, new network infrastructures which take into account the availability of optical fiber during the design could improve the network's performance.

1.2 A Truly Inconvenient Truth

In 1990 the World Wide Web (WWW) was invented, a service running over the Internet, allowing anyone with a connection to the world-wide network to access information from across the globe within seconds. Around this time, the consumer market started offering truly affordable Personal Computers (PCs). Thanks to this evolution, households everywhere soon had access to the WWW and its growth continues to this day.
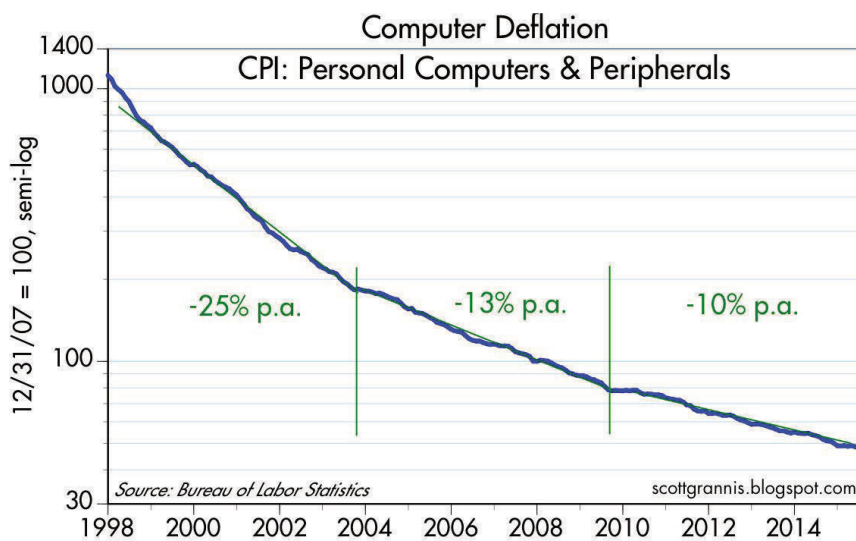
Figure 1.6: Personal Computer Price Evolution (1998 - 2015) [5]

Over the years PCs have become more affordable than ever and new technologies have enabled the development of portable devices such as laptops, tablets and smartphones. All of these are devices with impressive computational power for a fraction of the price of a PC in the early days. However, today the computational power of these devices is not the main interest of their owners. Owing to the widespread availability of Internet connectivity, these devices have become communication tools rather than computational ones. The evolution of electronics prices (Figure 1.6) combined with manufacturers truly flooding the market with devices (Figure 1.7) has gradually changed the *one-device-per-household* to a *multiple-devices-per-person* scenario. With communication devices penetrating every aspect of our everyday lives, it would be ignorant to deny this has a significant impact on our society.

Computers, smartphones, and tablet sales: 1975-2011

Unit sales (thousands)
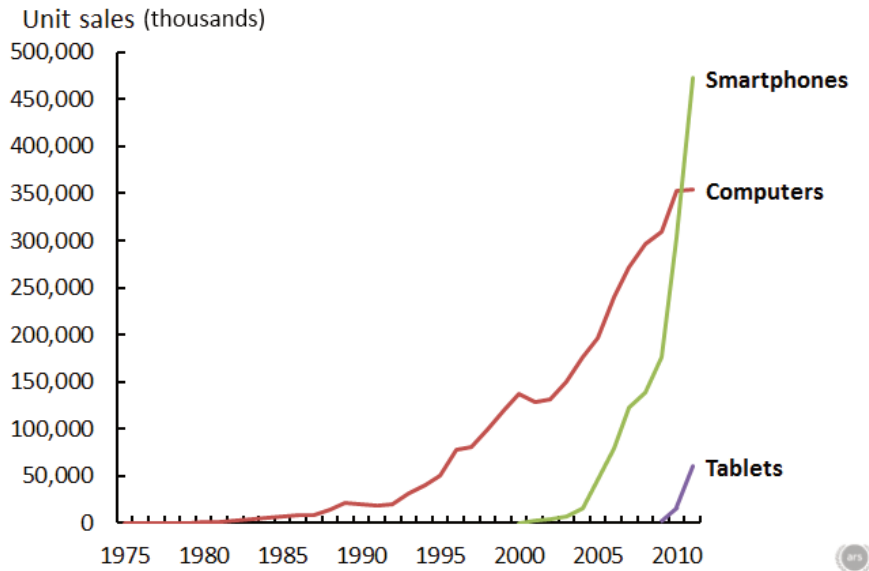
Figure 1.7: Computers, smartphones and tablet sales evolution [6]

1.2.1   Ecological and Economical Impact

It is often hard to predict the impact of a new technology, while it might have far-reaching consequences. In the case of the Internet numerous obvious advantages were associated with rolling out such as network: reduced traveling, worldwide knowledge and resources sharing, ... As a result, the Information and Communication Technology (ICT) sector has exhibited a rapid growth, offering a multitude of benefits to its adopters, but resulting at the same time in an increasingly large impact on the world.

For example, the global power consumption of communication networks in 2012 has been estimated at 350 TWh [7], or a staggering 1.8% of the total global power consumption at that time. As long as the Internet was of modest size, the power consumption to keep the network up and running was limited. However, with the network deployed today the power consumption has truly exploded and will continue to do so if no action is taken in the very near future.

The high power consumption associated with our Internet usage is an issue both from an ecological as from an economical point of view. The ecological impact is very hard to estimate, since some of the benefits of having the

Internet at our disposal actually have a positive impact on the global ecological footprint. However, with an estimated power consumption this high, one can hardly argue the ecological impact of our communication networks is not going to be significant.

Economically, one can easily understand the cost of energy becomes a major issue: at an estimated average cost of 8 eurocent/kWh, the energy to keep our communication networks turned on in 2012 was about 28 billion euro. You don't need to be an economical mastermind to see that a serious power consumption reduction is economically appealing.

Moreover, concern is rising regarding increased energy-dependency. With such vast amounts of energy needed to support our communication, any hiccup in energy production or distribution might be catastrophic for our economy. Furthermore, in [8] the author wonders if state-of-the-art ICT might become only accessible to those countries with nearly unlimited energy resources. Such a situation would result in an economic inequality recently described as the *digital divide*.

1.3 Towards sustainable communication

Over the past few years, multiple initiatives have been started in an effort to reduce the impact of communication networks on the environment, while still addressing the needs future networks have. In the context of this dissertation, two of these initiatives are highlighted here.

GreenTouch The GreenTouch consortium was founded in 2010 with the goal of reducing the power consumption of communication networks to a tiny fraction of what it is today. Realizing this goal basically means reinventing communication networks, which would require involvement of the diverse organizations throughout ICT. Therefore, a consortium was formed to bring all these organizations together.

DISCUS As part of the EUs Seventh Framework Programme (FP7), this project aims to exploit demonstrated technology and concepts needed to define and develop a radically new architectural concept that can enable an integrated wireless and Fiber-To-The-Premises (FTTP) future network which addresses the economic, energy consumption, capacity scaling, evolutionary, regulatory and service demand challenges arising from a FTTP enabled future. The key target of the DISCUS project is a cost-effective architecture for ubiquitous broadband services overcoming the digital divide.

1.4   Next-generation networks

Modern telecommunication networks are constructed as a three-level hier-
archical network (Figure 1.8), where each level is called a tier, not to be
confused with the ISP tiers mentioned earlier. Each tier can roughly be
identified with a geographical entity. The *core tier*, which typically uses a
mesh topology, is responsible for interconnecting continents and countries,
spanning distances ranging from hundreds to thousands of kilometers.
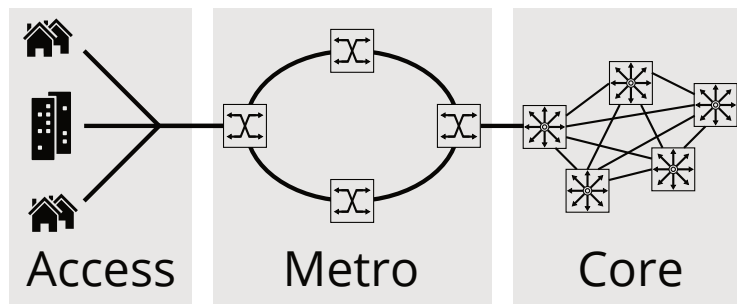
Figure 1.8: Modern telecommunication network hierarchy

The *metro tier* corresponds to a metropolitan area, and can roughly be seen
as the area a large city covers. Metro networks consist of a ring topol-
ogy interconnecting several Central Offices (COs) over tens to hundreds of
kilometers.

The *access tier* is the lowest level and provides connectivity to the end-
user. Access networks are designed to operate on distances of a few to
tens of kilometers. Contrary to the core and metro tier, the access tier is
deployed in a variety of configurations, such as bus, star or ring topologies.
Due to the hierarchical construction of a modern telecommunication net-
work it is clear the amount of network devices deployed in the access tiers
far outnumbers those in the core tiers (Table 1.1). Therefore, small power
consumption reductions in access network devices potentially have a much
bigger impact on the complete system than high power consumption reduc-
tions in power-hungry devices in the core network.

1.4.1   Advances in the access network

Contrary to the core and metro networks, where the interconnection medium
is optical fiber, the access network historically uses copper cable: the twisted

	Power consumption (W/device)	# of devices	BAU (GWh/yr)
Home	10	17,500,000	1,533
Access	1,280	27,344	307
Metro	6,000	1,750	92
Core	10,000	175	15

Table 1.1: 2015-2020 Italian Network Forecast: Device Density and BAU Energy Requirements [9]

pair telephone cable network and the coaxial cable CATV networks were readily available when the access networks had to be deployed. The Capital Expenditure (CapEx) costs associated with installing fiber are very high, and it would not have been sensible to ignore readily deployed networks with, at that time, more than enough bandwidth available.

1.4.1.1 Copper limitations

The legacy networks were installed with a completely different use-case in mind, and therefore used low-cost copper cables. Therefore, the copper access networks have serious impairments: not only do they have a limited bandwidth, they suffer from high losses, dispersion problems and crosstalk issues.
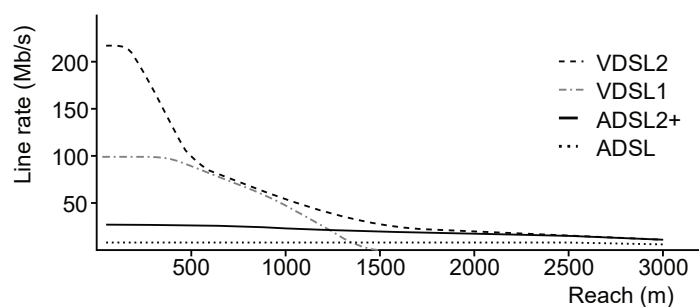
Figure 1.9: Rate vs Reach for different xDSL Technologies [10]

However, thanks to the development of several impressive advanced transmission techniques, service providers have steadily been able to provide

their customers with higher data rates, despite the low quality copper cables they are using. Unfortunately, these advanced techniques can only support limited distances. This is illustrated in Figure 1.9, where different Digital Subscriber Line (DSL) technologies are displayed with their maximum supported data rate versus distance between transmitter and receiver. The original Asymmetric Digital Subscriber Line (ADSL) could support up to 80% of the maximum data rate over distances longer than 3 km. On the other hand, Very-high-bit-rate Digital Subscriber Line 2 (VDSL2) supports a much higher data rate, but cannot support this data rate on distances over 200m. The emerging G.fast standard supports data rates up to 1.1 Gbit/s over a distance of 70m, but this quickly drops to 200 Mbit/s at distances of 200m. Although impressive achievements, this is not a sustainable way to cope with the ever-increasing demand for higher bandwidth. Moreover, these advanced transmission techniques all require extensive signal processing, making these solutions very power hungry.

1.4.1.2 All-optical access networks

At the time the technological advancements made a commercial all-optical access network possible, the cost of such a network was way too high to even consider. Today, the situation has changed: optical fiber offers a multitude of advantages over copper cable, as summarized in Table 1.2. Therefore, a Fiber-To-The-Home (FTTH) infrastructure is desirable, where optic fiber is distributed to every subscriber's premise [11]. Despite all technological advantages, converting a legacy copper access network to an all-optical access network requires a dazzling investment, which explains why this topology has not been globally deployed at this time.

However, the limited distances supported by the currently used advanced transmission techniques have already forced service providers to adopt an intermediate solution: Fiber-To-The-Curb (FTTC). FTTC is a hybrid fiber access solution which consists of routing fiber to a cabinet in the street, close to the end-user, without requiring to install fiber to every single end-user. In an FTTC scenario, the connection from the cabinet to the end-user, which is still copper, is known as the *last mile*. Converting this last mile to optic fiber is very expensive, owing to the cost of the required civil works. Although this is quite the investment, more and more service providers start to recognize it would turn their current situation of bandwidth scarcity to one of bandwidth abundance, which enables long term growth and creates potential for additional services on the network.

Property	Advantage
Size	The total diameter of an optical fiber (core, cladding and protection jacket) measures about $400\,\mu$m, a significant reduction from the 6mm diameter of coaxial cable. This is advantageous in cramped conduits in buildings and underground layout.
Weight	Due to the mass density difference and the smaller size, optical fiber yields a 10 to 30 % weight reduction compared to copper cable.
Bandwidth	Fiber has very high bandwidth, supporting data rates over 100 Tbit/s across one single standard SMF, as experimentally proven in [12].
Loss	Optical fiber has an attenuation less than $0.2\,$dB/km at $1550\,$nm, enabling transmission over several tens of kilometers without amplification.
Electrical interference	Since light is used, electromagnetic fields have no influence on transmission making it ideal in environments where strong electromagnetic fields are present.
Crosstalk	Very little light escapes an optical fiber, resulting in very good crosstalk characteristics.
Dispersion	At $1550\,$nm the dispersion of optical fiber is about $17\,$ps/(nm.km), while at $1300\,$nm it can be reduced to 0.
Environmental	Flammable or explosive environments pose no issue, since no sparks are ever generated by optical fiber.
Material availability	While copper is mined, and is scarce, silica is composed of oxygen and silicon, both available abundantly.
Multiplexing	A single fiber supports multiplexing of many wavelengths, increasing the potential datarate.

Table 1.2: Advantages of optical fiber [10]

1.4.2   Passive Optical Networks

One of the most cost effective and energy efficient optical access networks are Passive Optical Networks (PONs). PONs provide a connection between one Optical Line Terminal (OLT) and several Optical Network Units (ONUs), as displayed in Figure 1.10. Each ONU serves a number of end-users, depending on the configuration. A PON spans about 20 km, from the Central Office over a Remote Node (RN) to an ONU which in turn delivers the network traffic to the subscribers. As shown in Figure 1.10, the part of the network between the CO and the RN is called the feeder section, and consists of a single fiber. In a PON configuration, the RN is a passive optical power splitter, removing the need for power and drastically reducing maintenance at the RN. The distribution section is the part of the network starting from the RN to the different ONUs, effectively distributing the traffic to all subscribers.
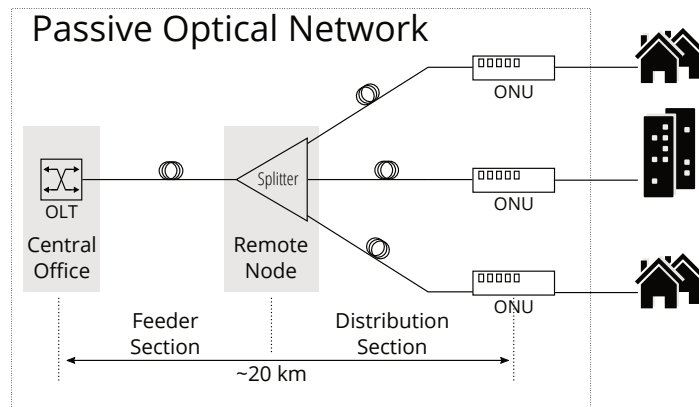
Figure 1.10: Passive Optical Network

1.4.3   The metro-access convergence - Long-Reach PONs

The evolution to all-optical access networks has suddenly changed the context of communication networks: there is no longer a physical difference between the transmission media used in the metro and in the access network. Therefore, in next-generation networks, any distinction between metro and access networks will be an artificial one.

By letting go of the legacy metro versus access network partitioning, the concept of the Long-Reach PON (LR-PON) was conceived. A LR-PON combines what used to be the metro and access network and treats them as
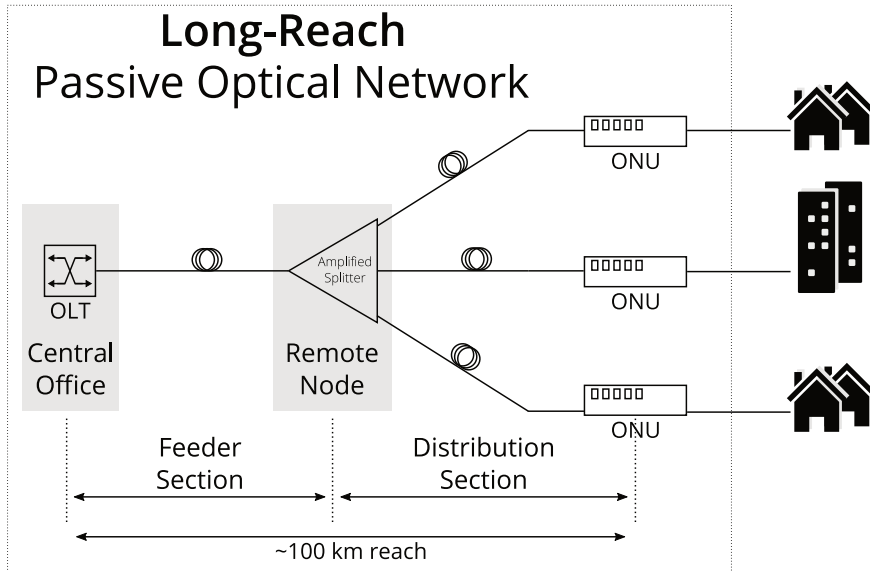
Figure 1.11: Long-Reach Passive Optical Network

one single network, spanning much larger distances (up to about 100 km) than a traditional Passive Optical Network (PON). The main architecture of a LR-PON is similar to a classical PON, consisting of two sections: (1) a feeder section and (2) a distribution section, as shown in Figure 1.11. However, the combined feeder and distribution section span a distance of around 100 km, effectively supporting a much longer reach, as the name suggests. The RNs are typically installed in old CO locations that have been consolidated to the LR-PON CO. Feasibility of such networks was already proven in 1998 by the ACTS PLANET lab demonstrator which was a technical implementation of SuperPON [13].

1.5   Challenges of next-generation networks

As technology advances, the expectations of the public advance with it: where in the early 1990s people were impressed by an Integrated Services Digital Network (ISDN) line offering up to 128 kbit/s, today people no longer consider this a workable connection.

Internet data traffic has been increasing steadily over the past few years, and is predicted to continue increasing in the foreseeable future, evident from the forecast shown in Figure 1.12. However, the ubiquitousness of
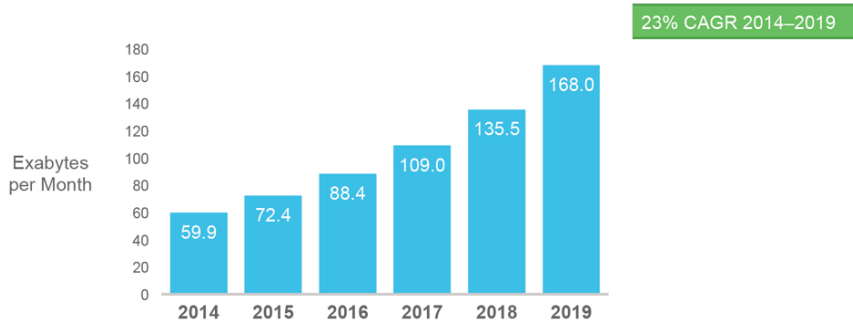
Figure 1.12: IP Data traffic Forecast [14]

the Internet, partly to be attributed to the strong uprising of smartphones, is rapidly changing the way people and industries see and use the Internet. This is the driving factor responsible for the emergence of new technologies such as Cloud Services and the Internet of Things, which are predicted to experience an explosive growth in the coming years (Figure 1.13).
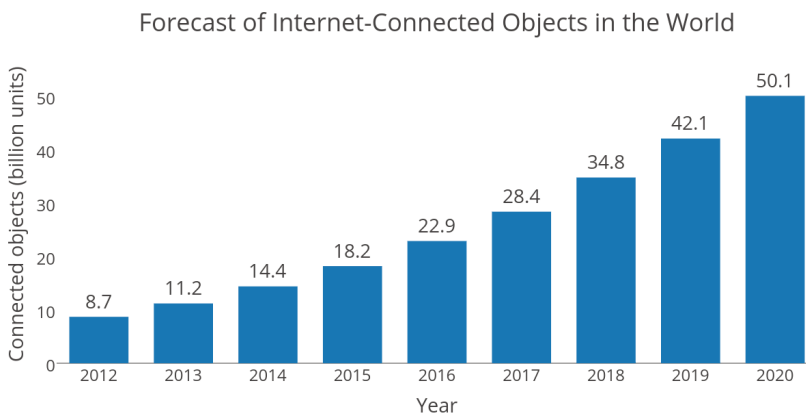
Figure 1.13: Forecast of Internet-Connected Devices in 2013 [15]

These new technologies have a strongly deviating data traffic profile compared to what is in use today. This will translate to a different set of requirements for the underlying communication network. For example, when the Internet of Things is in place, the Internet will be flooded with numerous tiny packets of data, a scenario which has not been taken into account when the currently deployed communication networks were designed. On the other hand, the tremendous increase in streaming services all require a low

packet-delay variation to offer a quality of service to the end-user. Again, this has not been incorporated in the design process of the currently deployed communication networks which are dominated by IP traffic, since such use of the network was not expected.

From these observations, we can conclude there are two main challenges to be addressed: (1) users desire higher bandwidths and (2) new services require more flexibility regarding bandwidth allocation and reduced latency. As previously mentioned, these challenges are complemented by the stringent need for power consumption reductions.

1.6 Overview of the work

This dissertation is based on the research the author has conducted over the last 5 years at the INTEC Design laboratory of the Department of Information Technology (INTEC) at Ghent University.

One of the main goals of the INTEC Design laboratory is to educate young engineers in advanced electronics by offering them the chance to start a doctoral training. Through close collaboration with industrial partners, as well as academic institutes, projects always have a strong applied nature. This requires the designer to take every step in the design cycle: from deriving specifications, designing the necessary circuits and test boards, to thoroughly testing the developed hardware. The main research topics handled by the Design group are RF and broadband communications, including fiber optics. The designs are highly innovative and often too complex to be tackled by an individual. Therefore, teamwork is a cornerstone of the projects within the Design Group. This allows the Design group to maintain its strong international reputation within the world of high speed electronics and fiber communication.

The previous sections of this chapter have introduced the reader to the 3 main observations that have defined the context wherein the research leading to this dissertation has been conducted. Firstly, there is the growing importance of communication networks power consumption, which can no longer be ignored. Therefore, next-generation networks will have to engage low power solutions. Secondly, the tremendous increase in data traffic can only be sustained by communication networks supporting higher line rates. Thirdly, the emergence of new technologies introduces a mixture of connected devices that have strongly deviating data traffic profiles from what our communication networks have seen in the past. As a result, the need for higher bandwidth flexibility will become apparent.

In this dissertation the Cascaded Bit-interleaving PON (CBi-PON) is presented as an answer to the various challenges next-generation networks are facing. CBi-PON offers a low-cost, low-power solution to the metro-access convergence problem while still supporting a very flexible dynamic bandwidth allocation scheme.

The author was part of the team that developed CBi-PON. It was the result of a collaboration between imec and Bell Labs in the context of the GreenTouch consortium. Furthermore, within the EU FP7 DISCUS project, the author was responsible for the development of a ONU Application-Specific Integrated Circuit (ASIC) supporting the goals of the DISCUS project. Moreover, the research efforts and results of the INTEC Design group were recognized through to the GreenTouch $1000\times$ award, as the INTEC Design group was one of the core partners of the Bit-interleaving PON (BiPON) and Cascaded BiPON team together with Bell Labs/Alcatel-Lucent and Orange Labs.

The major contributions of the author are to be found in the design, implementation and verification of the CABINET ASIC. In terms of design and implementation, the author was responsible in particular for several building blocks in the Analog Front- and Back-End and for the physical implementation of the MAC Preprocessor. Furthermore, contributions were made to the development of the CBi Protocol in collaboration with the team from Bell Labs/Alcatel-Lucent, from whom the idea of the CBi-PON originated.

1.7   Organization of this dissertation

In this chapter a brief history of the Internet has been given, revealing the background to which the currently deployed communication networks have gradually grown to their current state. Consequently, the issues regarding the power consumption of these communication networks have been highlighted, as well as some of the initiatives that have been taken by industry and governments to turn the tide. Furthermore, the change in needs and the accompanying challenges for next-generation metro/access networks due to the emergence of new technologies have shortly been discussed. In chapter 2, a first important step in dealing with the issues and challenges listed here is introduced: the Bit-interleaving PON (BiPON) protocol. Chapter 3 extends the ideas and concepts of the BiPON into a multi-level scenario, named the Cascaded Bit-interleaving PON (CBi-PON), in an effort to further reduce communication networks power consumption, while still addressing the challenges posed by new technologies. Subsequently, chapter

4 continues by discussing in detail the multi-rate, multi-mode ASIC implementation of a generic CBi device. In chapter 5 estimations on the power consumption reduction are presented, showing the potential of the concept in real-world scenarios. Finally, chapter 6 concludes the dissertation by providing a summary of the most important results and discussing opportunities for future research.

References

[1] Peter H Salus and G Vinton. *Casting the Net: From ARPANET to Internet and Beyond...* Addison-Wesley Longman Publishing Co., Inc., 1995.

[2] Merit Network Inc. NSFNET 56K Backbone Map, July 1986 - July 1988, 2011.

[3] Merit Network Inc. NSFNET T1 Backbone Map 1991, 2011.

[4] Merit Network Inc. NSFNET T3 Backbone Map Fall 1992, 2011.

[5] Scott Grannis. Computer Deflation - CPI: Personal Computers & Peripherals, 2015.

[6] Jeremy Reimer. From Altair to iPad: 35 years of personal computer market share, 2012.

[7] S. Lambert, W. Van Heddeghem, W. Vereecken, B. Lannoo, D. Colle, and M. Pickavet. Worldwide electricity consumption of communication networks. *Optics express*, 20(26):B513–B524, 2012.

[8] Maruti Gupta and Suresh Singh. Greening of the Internet. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 19–26. ACM, 2003.

[9] Raffaele Bolla, Franco Davoli, Roberto Bruschi, Ken Christensen, Flavio Cucchietti, and Suresh Singh. The potential impact of green technologies in next-generation wireline networks: Is there room for energy saving optimization? *Communications Magazine, IEEE*, 49(8):80–86, 2011.

[10] Renato Vaernewyck. *High-speed low-power modulator driver arrays for medium-reach optical networks*. PhD thesis, Ghent University, 2014.

[11] Eileen Connolly Bull. FTTH Handbook, Edition 6. *Fiber To The Home Council Europe*, 6, 2014.

[12] Dayou Qian, Ming-Fang Huang, Ezra Ip, Yue-Kai Huang, Yin Shao, Junqiang Hu, and Ting Wang. High capacity/spectral efficiency 101.7-Tb/s WDM transmission using PDM-128QAM-OFDM over 165-km SSMF within C-and L-bands. *Journal of Lightwave Technology*, 30(10):1540–1548, 2012.

[13] Jan Vandewege, Xing-Zhi Qiu, Brecht Stubbe, Chris Coene, Peter Vaes, Wei Li, Jan Codenie, Claire Martin, H Slabbinck, Ingrid Van de Voorde, et al. A lab demonstration of a SuperPON optical access network. In *Broadband Communications*, pages 69–80. Springer, 1998.

[14] Cisco Visual Networking Index. Forecast and Methodology, 2014-2019 White Paper. Technical report, Technical Report, Cisco, 2015.

[15] Cisco Consulting Services. The Internet of Everything (IoE) - Connections Counter, 2013.

2

Bit-Interleaving PON

The research presented in this dissertation introduces the Cascaded Bit-Interleaving PON, a network architecture concept which is thoroughly discussed in Chapter 3. The Cascaded Bit-interleaving PON makes extensive use of the Bit-interleaving PON (BiPON). Since BiPON is relatively unknown, but of major importance for this work, this chapter attempts to briefly introduce the reader to the concepts and particularities making the Bit-Interleaving PON a superior solution in comparison to traditional PON protocols.

2.1   The need for BiPON

As extensively covered in Chapter 1, the dramatically increasing power consumption of communication networks is becoming a serious issue. Therefore, the GreenTouch consortium, founded in 2010, set the goal of presenting a solution for communication networks by 2015, wherein the power consumption is reduced with a factor $1000\times$ compared to a reference architecture defined by the consortium in 2010.

To achieve such a dazzling power reduction, different work groups were formed, among which the Wireline Access Networking group. The fixed access network contains a stunning 94% of all devices in the network, and is consequently responsible for around 70% of the total networks power

consumption [1]. This large share in the total power consumption shows the biggest improvement is to be found in the access network.

Moreover, the active devices in a PON are Optical Line Terminals (OLTs) which typically servce 32 up to 128 ONUs. While an OLT generally uses around 500 mW per ONU, the ONU itself consumes around 6 W [2], which means the ONUs are clearly the major contributors to the access networks power consumption. Therefore, the largest power reduction is to be expected when concentrating on the ONUs. The following section gives an overview of the traditional PON protocols, introducing the reader to the context wherein BiPON was developed. Subsequently, Section 2.3 clarifies how BiPON differentiates from these traditional protocols and explains the mechanisms that form the basis for the impressive power consumption reductions in the ONU that have been demonstrated.

2.2 An overview of PON protocols

This section provides the reader with an overview of the different PON protocols and the evolution that has taken place. This allows to fully understand the context in which the BiPON protocol was developed and the advantages it brings.

2.2.1 Multiplexing schemes

The PON architecture accommodates multiple users on the same shared physical medium: the optical fiber. However, to share this fiber, a mechanism must be used to allow multiple data streams.

Compare this with a multi-user conversation: when you're in a group and several people ask the same person a different question, this person will have to find a way to answer all of them. However, this is not be possible at the same time, in the same way. He could answer them one at a time, or he could answer simultaneously by answering one by speech and another in writing, assuming he is capable of doing that.

The mechanism to cope with this shared medium is what we call *data multiplexing*, and currently deployed PONs are mostly Time Division Multiplexing (TDM) PONs. Time Division Multiplexing organizes the multiple access of the shared medium by assigning all receivers a designated time slot.

As shown in Figure 2.1, the OLT broadcasts all downstream data to the ONUs. Upon reception, each ONU extracts the packets addressed to it and

discards all other packets. In the upstream direction, the OLT dynamically allocates specific transmit time slots to each ONU. By means of a *grant packet* the ONU is informed of its transmit time slot, during which it is allowed to send data to the OLT. The OLT schedules these time slots such that no collisions occur. It is clear this method requires a precise synchronization of transmission time between all ONUs and the OLT. However, since TDM was a technique already deployed in copper, methods and protocols to achieve this synchronization were readily available.
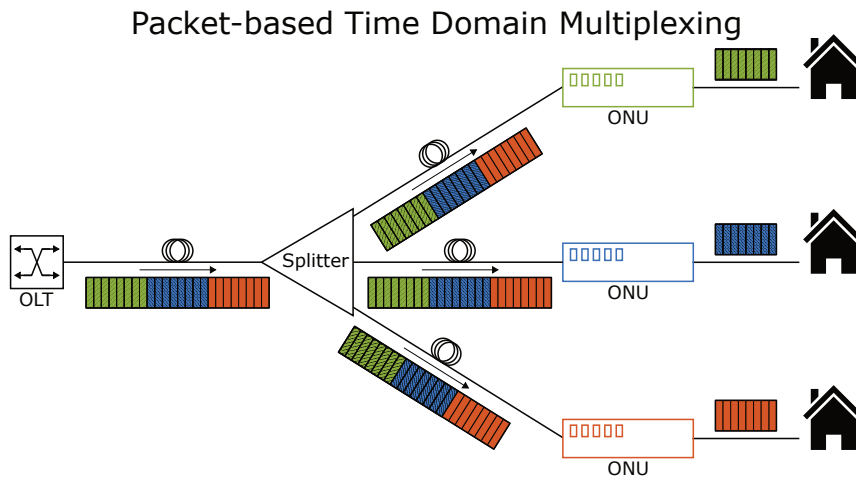
Figure 2.1: Packet-based Time Domain Multiplexing

Alternative data multiplexing schemes are Wavelength Division Multiplexing (WDM) and Orthogonal Frequency Division Multiplexing (OFDM), possibly offering higher data rates than TDM. Although these techniques are getting a lot of research attention, the cost of deployment is still quite high, and massive deployment has not started so far.

It is important to note that these schemes are not necessarily mutually exclusive. For example, one can combine TDM and WDM as TWDM-PON, which offers advantages of both techniques regarding cost and performance. In fact, TWDM-PON has been selected by the Full Service Access Network (FSAN) and the International Communication Unit (ITU) as the standard multiple access scheme for Next-Generation PON 2 (NG-PON2).

2.2.2   PON Standards

PON standardization began in 1995, when the FSAN initiative was formed. FSAN is a working group formed by network operators and telecom vendors in an effort to lower the costs of optical network equipment by introducing standards for the industry stimulating volume production. An overview of the PON standards is given in Table 2.1, while the timeline of their introduction is shown in Figure 2.2.
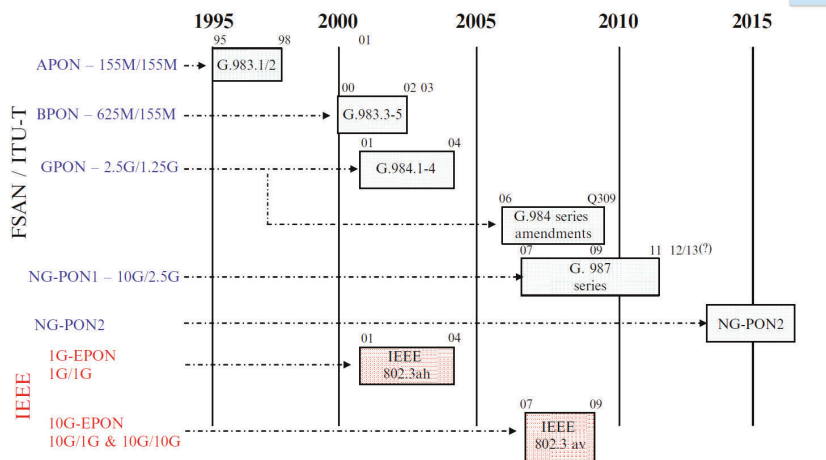
Figure 2.2: PON Standardization Timeline [5]

The first PON standard developed by FSAN and ITU was Asynchronous Transfer Mode (ATM) PON (APON). APON, standardized in 1998, was connection-oriented and strongly driven by voice-type communication, which dominated network traffic at that time. However, network traffic quickly changed and the Broadband PON (BPON) standard (ITU-T G.983), an amendment on APON, was finalized in 2002.

However, the underlying ATM PON proved to be an inefficient solution to transfer the data on the network, which was mostly IP based. To deal with this, both the ITU and the Institute of Electrical and Electronics Engineers (IEEE) suggested a new standard, while also increasing bandwidths, to cope with the rising network traffic.

The IEEE founded the Ethernet in the First Mile (EFM) initiative, which proposed a PON standard based on Ethernet: suitingly called Ethernet PON (EPON) or GE-PON. This was standardized in 2004 as IEEE 802.3ah. EPON modified the IEEE 802.3 standard to support the Point-to-MultiPoint

	BPON	EPON	GPON	10G EPON	XG-PON	NG-PON2
Organization	ITU-T	IEEE	ITU-T	IEEE	ITU-T	ITU-T
Standard	G.983	802.3ah	G.984	802.3.av	G.987	G.989
Downstream λ	1490 nm	1490 nm	1490 nm	1577 nm	1577 nm	1596-1603 nm
Upstream λ	1310 nm	1310 nm	1310 nm	1270/1310 nm	1270 nm	1524-1544 nm
Typical DS rate	622.08 Mbit/s	1 Gbit/s	2.488 Gbit/s	10 Gbit/s	9.953 Gbit/s	$4\times$ 9.953 Gbit/s
Typical US rate	155.52 Mbit/s	1 Gbit/s	1.244 Gbit/s	1 or 10 Gbit/s	2.488 Gbit/s	$4\times$ 2.488 Gbit/s

Table 2.1: PON Standards [3, 4]

(P2MP) connectivity which is characteristic to PON. This approach clearly demonstrates the view of the IEEE on IP as the dominant network technology.

ITU, on the other hand, did not share this view. The members of the ITU had large stakes in different services: ATM for voice, Ethernet for data and TDM-traffic. This required the new standard to natively support these services, which was achieved in the proposal for the Gigabit-capable PON (G-PON), standardized as G.984.
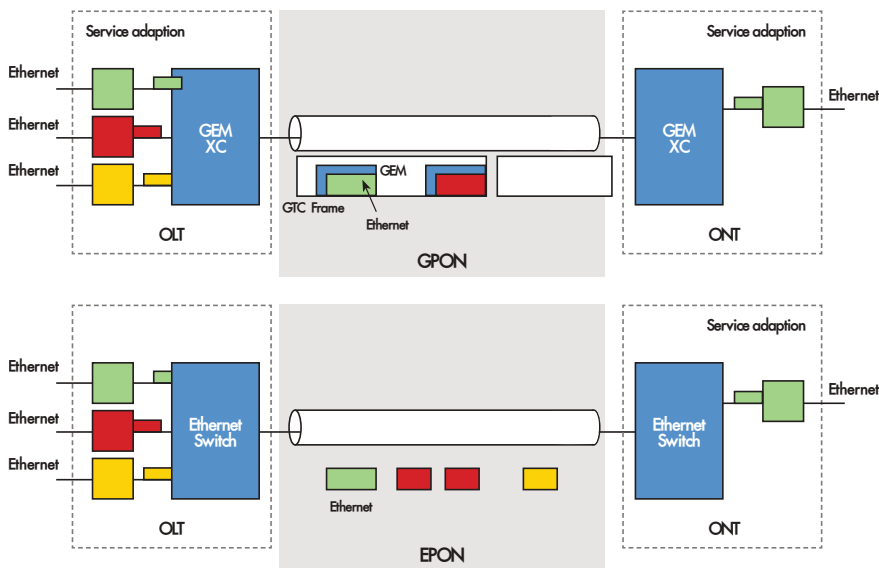
Figure 2.3: Generic Framing Procedure (GFP) in GPON and EPON [6]

G-PON got rid of the underlying ATM, but introduced the Generic Framing Procedure (GFP). Ethernet and TDM-traffic is encapsulated by the Generic Encapsulation Method and is combined with ATM packets and *Operation, Administration and Maintenance* (OAM) data on the G-PON Transmission Convergence (GTC) layer. The Generic Framing Procedure for both GPON and EPON is outlined in Figure 2.3.

It is clear G-PON does not support native Ethernet traffic and adds an extra layer of complexity compared to EPON (Figure 2.4). However, the long guard times specified for EPON reduce the transmission efficiency, which makes G-PON, despite its lack of native Ethernet support, a more appealing solution.

Today, we see massive deployment of EPON in Japan and South Korea,
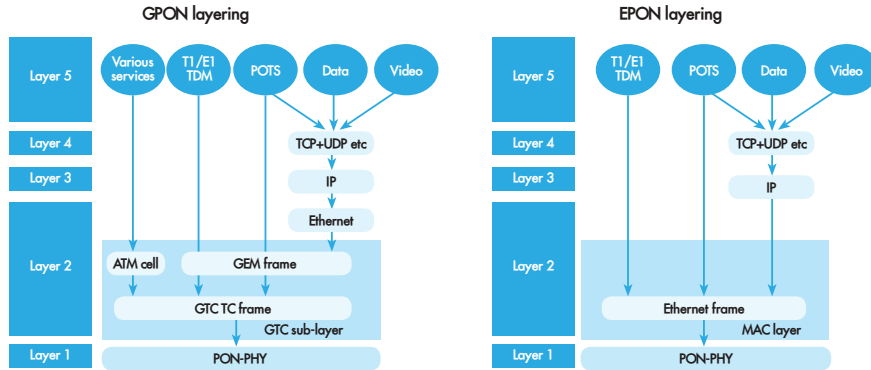
Figure 2.4: GPON vs EPON Protocol [6]

while GPON is the dominant technology in North America. China recently embarked on a large-scale deployment using a mix of EPON and GPON. This shows that it is unlikely that one of the two standards will ever dominate, due to the large installed user base of both standards.

2.2.3 Next-generation PONs: NG-PON

Even with these gigabit-capable PONs available, in the near future even higher bandwidths will be required. New emerging services like next-generation streaming 3D HDTV, remote medical monitoring or online gaming will require much higher bandwidths.

After the development of G-PON, FSAN and ITU-T started work on next-generation PONs (NG-PONs). The introduction of NG-PONs will, in the view of FSAN and ITU-T, happen in two phases: NG-PON1 and NG-PON2. NG-PON1 is supposed to be the mid-term upgrade of the PON network, while NG-PON2 is regarded as the long-term evolution, as shown in Figure 2.5.

2.2.3.1 NG-PON1

NG-PON1 technology is seen as the *evolutionary growth* phase of the NG-PON transition. Therefore, its goals are very clear: provide higher performance in terms of bandwidth using the existing Optical Distribution Network (ODN), while co-existing with legacy PONs.

With respect to these requirements, the ITU developed 10-Gigabit-capable PON (XG-PON) (G.987) as an extension on G-PON, while the IEEE ex-
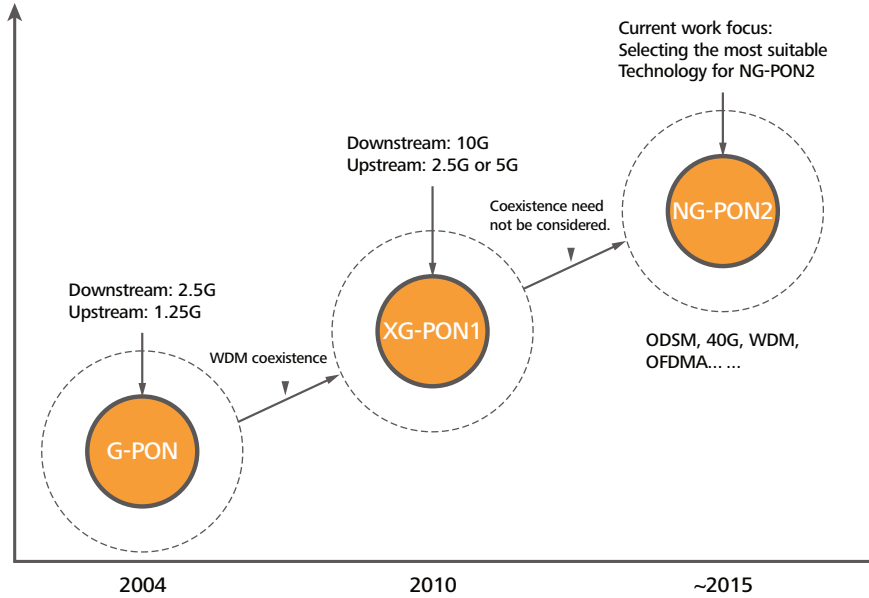
Figure 2.5: FSAN Next-Generation PON Roadmap in 2010 [7]

tended EPON to 10 Gbit/s Ethernet PON (10GE-PON) , IEEE 802.3av. Both offer 10 Gbit/s downstream rates and 2.5 Gbit/s upstream rates. However, 10GE-PON also defines a symmetric mode where 10 Gbit/s upstream data rate is supported.

2.2.3.2 NG-PON2

NG-PON2 takes a very different approach: NG-PON2 should foster the *revolutionary change*. While a main concern for NG-PON1 is backwards compatibility, i.e. co-existence with legacy PONs on the existing ODNs, this is no longer true for NG-PON2. Although NG-PON2 is coined as being disruptive, extensive re-use of the existing ODN is desired to be cost-effective.

NG-PON2 was standardized in 2015, offering a total network throughput of 40 Gbit/s, which corresponds to up to 10 Gbit/s for each subscriber, up- and downstream. NG-PON2 uses Time and Wavelength Division Multiplexing (TWDM) as access technology for both upstream and downstream directions, and is therefore called a TWDM-PON. A TWDM-PON is basically a TDM-PON with the added dimension of WDM. In the case of NG-PON2, 4 wavelengths have been selected for the WDM. Even though co-existence with older (e.g. GPON) PON technology was no requirement,

the wavelength selection shows that the standard nevertheless supports this (Figure 2.6).
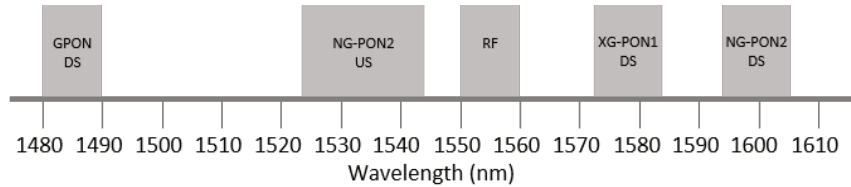
Figure 2.6: Optical network spectrum including NG-PON2 wavelengths [8]

This wavelength selection strategy ensures re-use of the original fiber, but OLTs and ONUs have to be replaced to support the wavelength-division multiplexing: the OLT needs a wavelength multiplexer, while the ONUs need to be accommodated with an active tunable filter to select the appropriate wavelength. In case a static wavelength allocation strategy is used, it suffices to populate the ONUs with a static filter.

One could argue that NG-PON2 was a chance to be disruptive, an opportunity to get rid of legacy inefficiencies, but the final decision did not take advantage of it and is rather conservative.

2.2.4 Power consumption in PONs

Currently installed PONs mostly use a time division multiplexing access scheme, since all standardized PON protocols are TDM-based. The ongoing demand for higher bandwidths has reflected on the TDM-based PON protocols by making them support higher and higher line rates.

However, the dynamic power consumption of CMOS circuits used in the ONUs scales proportionally with the operating frequency. Moving towards higher line rates results in higher operating frequencies, and therefore a higher dynamic power consumption. This problem is partly alleviated by the never-ending scaling of CMOS processes: the use of deep sub-micron CMOS technology reduces dynamic power consumption, though static power consumption rises.

2.2.4.1 Energy saving

Since the negative impact of the increasing power consumption is becoming more clear every day, both the ITU-T and IEEE have started initiatives

to focus on energy saving in optical networking. The ITU-T developed the G.Sup45 recommendation for power conservation in GPON, while the IEEE founded the IEEE 802.3az taskforce. Today, focus is mainly on sleep modes: ideally an ONU is only operating when receiving its own payload. However, in practice such an ideal scenario is not feasible, for example because there is a relatively long start-up time needed for the clock-and-data recovery circuits. This reduces the effectiveness of such additions.

2.2.4.2 Future

Although advances in CMOS technology promise substantial reductions in dynamic power consumption, the rising demand for higher bandwidths and its associated higher line rates will result in a net power consumption increase. To break this proportionality between power consumption and increasing bandwidths, a disruptive network architecture is required. The energy inefficiencies of older technologies were not significant when these technologies were designed, but with the rising bandwidths, these have become evident. Only by re-designing the network from the ground up we can escape from these inherited energy inefficiencies.

2.2.5 (X)G-PON power deficiency

The main focus of the Bit-Interleaving PON was to provide a power efficient alternative to XG-PON. This section will focus on revealing the power deficiency inherent to the XG-PON technology

As explained before, (X)G-PON is a TDM-PON using the GTP. Upon receiving a GTP frame at the full 10 Gbit/s line rate, the ONU has to perform a complex series of processing steps to extract the actual payload addressed to the ONU. Figure 2.7 presents a schematic overview of the receiver payload extraction, revealing the different steps that have to be taken.

It is evident from Figure 2.7 that a large part of the processing chain has to operate at the line rate, which is very power hungry. Only the very last part of the chain operates at the user rate, which is typically only 1 to 3% of the line rate and is therefore a lot less power hungry.

This clearly indicates the power inefficiency present in the processing scheme defined by XG-PON: although the user only needs data at the user rate, most of the processing is done at the much higher line rate.

10 Gbps From PON

OPTICAL/ELECTRICAL
FRONT END
(APD - TIA - LA)

CLOCK AND DATA
RECOVERY

DESERIALIZER

SYNC DETECTOR

WORD ALIGNER

DESCRAMBLER

FEC DECODER

XGTC/PLOAM RECEIVER

XGEM RECEIVER

100% of PON Traffic

REASSEMBLY

HEADER
PARSER

DECRYPTION

UNI MAC/PCS/PMD

1% to 3% of
PON Traffic

10/100/1000 Mbps To User
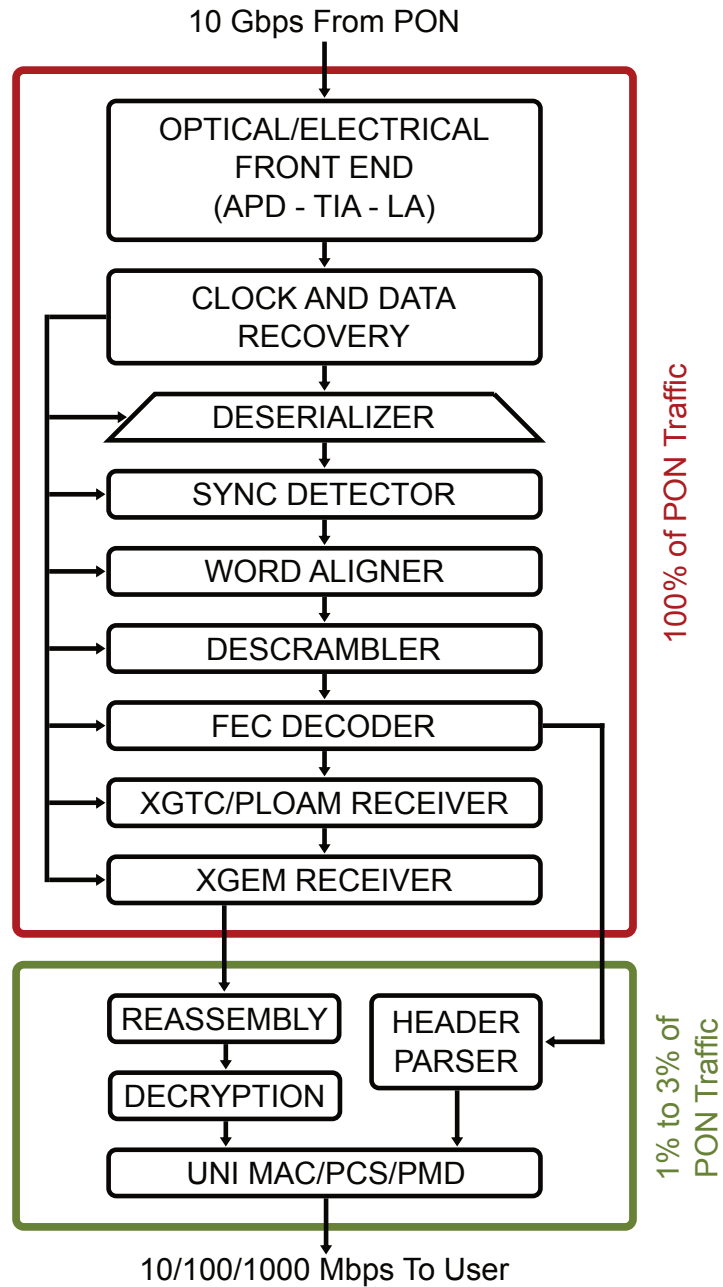
Figure 2.7: XG-PON Processing

2.3   Bit-interleaving PON

The bit-interleaving PON addresses the manifest power inefficiency present in the XG-PON protocol. The idea is to minimize the part of the receiver processing chain that needs to operate at the line rate.

2.3.1   Bit-based TDMA

The main differentiating feature of a Bit-Interleaving PON is the use of a bit-based Time Division Multiple Access (TDMA) instead of a packet-based TDMA (visualised in Figure 2.1). Bit-based TDMA is shown schematically in Figure 2.8. Instead of sending all bits for a single ONU all at once in a single packet, every next bit the OLT sends is addressed to a different ONU.
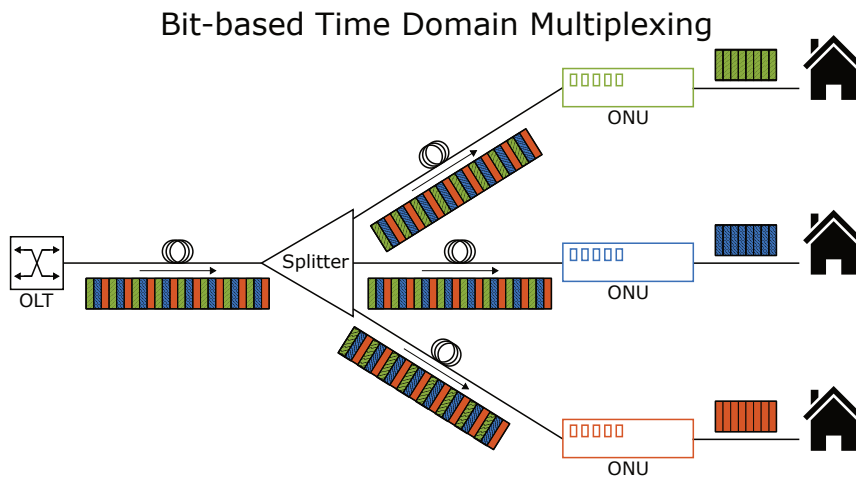
Figure 2.8: Bit-based TDMA

To fully leverage the potential of using a bit-based TDMA scheme, the proposed BiPON protocol modifies the transmission of a frame to achieve a bit-interleaving structure. This is presented in Figure 2.9: bits in a single row are intended for a single ONU, which means each row corresponds to a different ONU. Transmission, however, is done column by column, as indicated by the *transmission order* arrow on the figure. This transmission strategy provides the desired bit-interleaving arrangement.

Corresponding to an XG-PON frame, a BiPON frame has a fixed length of 125 μs. Furthermore, it consists of two sections: (1) the BiPON Header
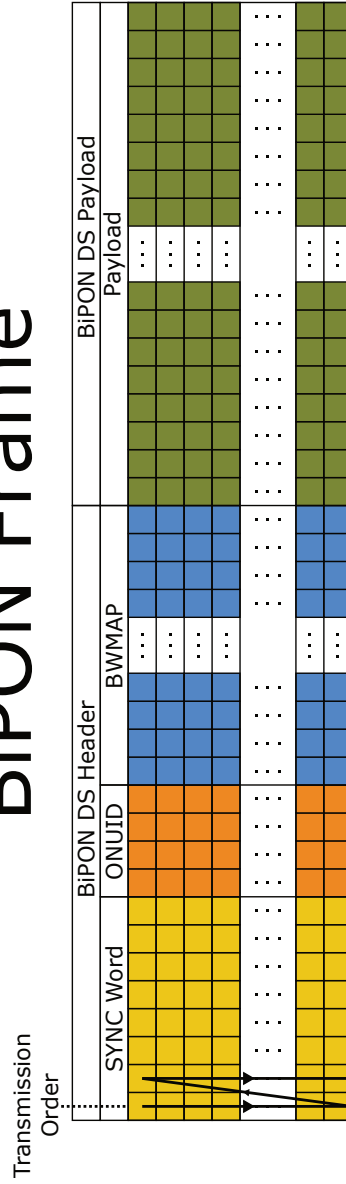
Figure 2.9: BiPON Frame Structure

and (2) the BiPON Payload. The BiPON Header is further divided in a SYNC Word, an ONUID and a Bandwidth Map (BWMAP) field. A BiPON frame accommodates a fixed number of 256 ONUs, each having a dedicated header channel.

The ONU determines its header channel by matching the SYNC word. Based on this match and the extracted ONU-ID, the offset for the header channel assigned to the ONU is easily calculated. This greatly reduces the required receiver complexity, since the decimation operation is simplified purely to a sub-sampling operation, as opposed to the much more complicated de-serialization and word alignment procedures used in XG-PON.

Furthermore, when comparing the processing chain to receive a BiPON frame (Figure 2.10) to the processing needed for XG-PON (Figure 2.7), it is obvious the operating frequency of the processing blocks is reduced to the user rate very early in the processing chain. Thanks to the bit-interleaving structure and the early down-sampling, all blocks following the decimator block can operate at the user rate, resulting in a tremendous power consumption reduction.

2.3.2 Dynamic Bandwidth Allocation

An important feature of the BiPON protocol is the availability of a flexible Dynamic Bandwidth Allocation (DBA) mechanism. While the header part of the frame is fixed-rate, reception of the payload section is subject to configuration by the bandwidth map BWMAP.

The bandwidth allocation information is transferred in the form of (S_i, K_i) doublets, with S_i the sampling rate and K_i the starting offset of ONU$_i$. This information unambiguously defines how the ONU should sample the payload section to extract the appropriate payload bits. Such an implementation of dynamic bandwidth allocation allows the OLT to alter ONU allocations on a BiPON frame-to-frame basis, with sub-frame resolution, based on the instantaneous traffic profile of the ONUs, making it extremely flexible.

2.3.3 Upstream traffic

BiPON focuses on the downstream traffic, and does not propose any adjustments for the upstream network. This is easily understood when taking into account the nature of the power deficiency in XG-PON, which BiPON is trying to eliminate. The root cause of the inefficiency is that all data is being processed, while a receiver does not use all data and discards the largest
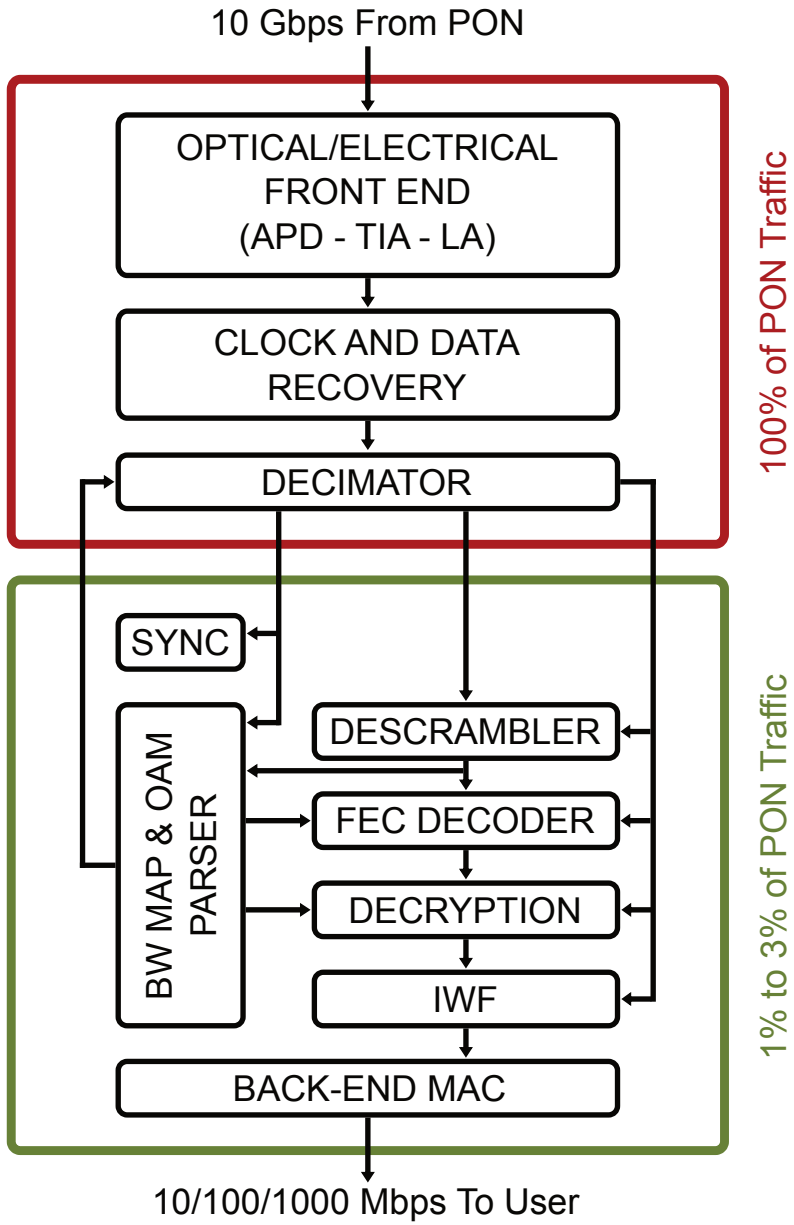
Figure 2.10: BiPON Processing

part. While this is true for the downstream direction, it is not the case for the upstream direction. In the upstream direction, the OLT needs to capture all data to retransmit it further upstream. There is no energy wasted with a discarding operation that can be eliminated, and therefore BiPON does not target the upstream direction.

2.4 Demonstrated Results

After introducing the basic principles and advantages of BiPON, this chapter is concluded by presenting the power consumption reductions achieved by the 10 Gbit/s BiPON implementation demonstrated in 2012.

2.4.1 FPGA-based Implementation

To demonstrate the concept of the Bit-Interleaving PON and to compare it with the XG-PON standard, a testboard was designed using an off-the-shelf XG-PON transceiver and a commercial 10 Gbit/s Clock-and-Data Recovery (CDR) sold by Vitesse. On an Altera Stratix IV Field-Programmable Gate Array (FPGA), both the XG-PON and the BiPON protocols were implemented to make a fair comparison. Further details on the implementation can be found in [9].

The results of the power consumption comparison are shown in Table 2.2. It is immediately clear the XG-PON protocol power consumption does not scale with the desired user rate, while the Bit-Interleaving PON scales very well. Moreover, for the same user rate, BiPON consumes 35× (@ 1.25 Gbit/s) to 180× (@ 10 Mbit/s) less than XG-PON. It has to be noted that these results only take the *dynamic* power consumption of the FPGA into account, not the *static* power consumption, nor the power consumption of other components.

2.4.2 ASIC Implementation

Since the FPGA implementation proves the potential of the Bit-Interleaving PON, an ASIC implementation was developed. Alongside the Medium Access Control (MAC) preprocessing for the BiPON protocol, a suitable CDR was implemented. Table 2.3 summarizes the dynamic power consumption measurements of the ASIC implementation and compares these against the XG-PON and BiPON FPGA implementation. It is important to note that, since the ASIC implementation incorporates a CDR, the power consumption of the commercial Vitesse CDR used in the FPGA demonstration has been added to the power numbers for the FPGA implementation.

Dynamic Power Consumption [W]	XG-PON FPGA	BiPON FPGA
Active @ 1.25 Gbit/s	3.6	0.102
Active @ 625 Mbit/s	3.6	0.065
Active @ 312 Mbit/s	3.6	0.044
Active @ 156 Mbit/s	3.6	0.037
Active @ 78 Mbit/s	3.6	0.029
Active @ 39 Mbit/s	3.6	0.024
Active @ 20 Mbit/s	3.6	0.022
Active @ 10 Mbit/s	3.6	0.020
Idle without sleep	3.6	0.009
Idle with sleep	2.6	0.009

Table 2.2: Estimated dynamic power consumption XG-PON vs BiPON on FPGA [9]

Dynamic Power Consumption [W]	XG-PON FPGA	BiPON FPGA	BiPON ASIC
Active @ 1.25 Gbit/s	3.698	0.200	N/A
Active @ 625 Mbit/s	3.698	0.163	0.105
Active @ 312 Mbit/s	3.698	0.142	0.087
Active @ 156 Mbit/s	3.698	0.135	0.078
Active @ 78 Mbit/s	3.698	0.127	0.064
Active @ 39 Mbit/s	3.698	0.122	0.060
Active @ 20 Mbit/s	3.698	0.120	0.055
Active @ 10 Mbit/s	3.698	0.118	0.050
Idle without sleep	3.698	0.107	0.050
Idle with sleep	2.698	0.107	0.039

Table 2.3: Estimated dynamic power consumption XG-PON vs BiPON on FPGA [9]

Compared to the BiPON FPGA implementation, the ASIC implementation shows $1.6\times$ (@ 625 Mbit/s) to $2.36\times$ (@ 10 Mbit/s) reduction.

2.5   Conclusion

In this chapter the need for the Bit-Interleaving PON was highlighted. After an introduction to the currently standardized PON protocols and the roadmap for future standardization, the power consumption in PONs was discussed. In this discussion, it was clarified how the current power-saving efforts are not future-proof. Subsequently the power deficiency that exists in today's PON standards was revealed.

The reader was introduced to bit-based time domain multiple access (TDMA), the disruptive paradigm enabling the bit-interleaving PON. The bit-interleaving PON was briefly introduced, stressing the advantages of the protocol over XG-PON.

Finally, the demonstrated results were summarized to show the potential of the technology that has formed the basis of the research in this work.

References

[1] Raffaele Bolla, Roberto Bruschi, Franco Davoli, and Flavio Cucchietti. Energy efficiency in the future internet: a survey of existing approaches and trends in energy-aware fixed network infrastructures. *Communications Surveys & Tutorials, IEEE*, 13(2):223–244, 2011.

[2] GreenTouch Foundation. Improving the energy efficiency of residential fixed access networks by more than $250\times$ by 2020. 2016. [Online].

[3] Tommaso Muciaccia, Fabio Gargano, and Vittorio Passaro. Passive Optical Access Networks: State of the Art and Future Evolution. In *Photonics*, volume 1, pages 323–346. Multidisciplinary Digital Publishing Institute, 2014.

[4] ITU-T Recommendation G.989.2 : 40-Gigabit-capable passive optical networks 2 (NG-PON2): Physical media dependent (PMD) layer specification. Technical report, ITU-T.

[5] FS.COM. Comparison of EPON and GPON: Evolution of PON, 2015.

[6] CommScope Solutions Marketing. GPON - EPON Comparison. White paper, CommScope, Inc., October 2013.

[7] Huawei. Next Generation PON Evolution. White paper, Huawei Technologies Co., Ltd., 2010.

[8] Salem Bindhaiq, Abu Sahmah M Supa, Nadiatulhuda Zulkifli, Abu Bakar Mohammad, Redhwan Q Shaddad, Mohamed A Elmagzoub, Ahmad Faisal, et al. Recent development on time and wavelength-division multiplexed passive optical network (TWDM-PON) for next-generation passive optical network stage 2 (NG-PON2). *Optical Switching and Networking*, 15:53–66, 2015.

[9] Christophe Van Praet. *Techniques to Reduce Energy Consumption in Next-Generation Access Networks*. PhD thesis, Ghent University, 2014.

3

Cascaded Bit-interleaving PON

While the Bit-interleaving PON showed very promising results regarding power consumption, the work does not end there. In the fast-changing landscape of communication networks, there is a never-ending need for lower power consumption and higher data rates. On top of that, data traffic will not only rise in the coming years, but will also take on new forms. The emergence of the Internet-of-Things and the increasingly widespread adoption of cloud services will most certainly have its impact on the data traffic profiles. To deal with this change, next-generation networks need to take this into account.

This chapter introduces the Cascaded Bit-interleaving PON, which was developed in an effort to meet the requirements of next-generation networks.

3.1   Metro-Access convergence & Long-Reach PON

Although BiPON is quite disruptive, it does preserve compatibility with the physical network: the Optical Distribution Network (ODN) used for an XG-PON network can be transformed to a BiPON network simply by replacing the electronics at the OLT side and the electronics at the ONU side. This means the core-metro-access hierarchy of communication networks is maintained.

However, in the past years this core-metro-access structure has been the

subject of discussion. Traffic profiles in the metro and in the access network are similar, and there is no obvious reason to make the distinction between these two networks. It is therefore predicted that the metro and access networks will cease to exist as separate entities, and will merge into a single metro-access network. Such a consolidation would allow further optimization of network resources.

In 2002 Davey and Payne showed that a Long-Reach PON (LR-PON) [1] enables significant cost savings, owing to the change in the network such that the electronic equipment and real estate required for the Central Office (CO) decreases considerably. Furthermore, the reach targeted in such a LR-PON spans both the metro and the access network, and is therefore a step towards the envisioned unified metro-access network.
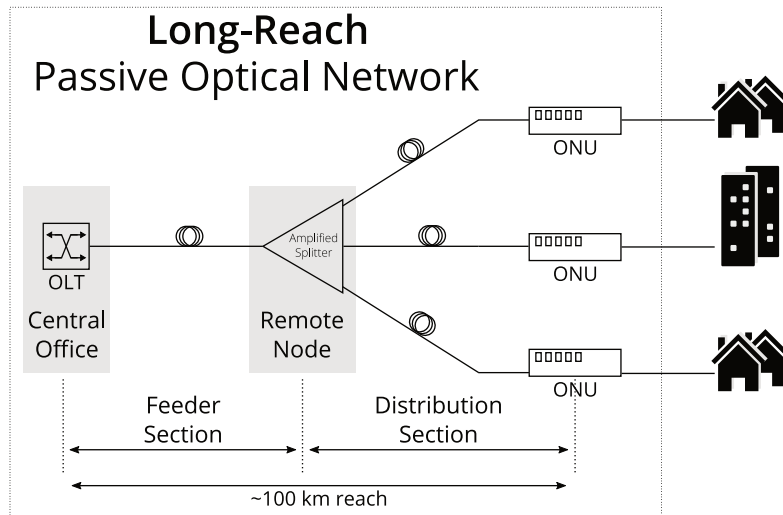
Figure 3.1: Long-Reach PON

The core idea of an LR-PON is to support a much larger split ratio and increase the reach of a PON, hence the *long-reach* designation. This can be achieved by replacing the passive splitters used in a traditional PON by amplified splitters, as shown in Figure 3.1. This addition to the network has ignited a discussion regarding the use of *PON* in the name, since the optical amplifiers in the splitters make the network active instead of passive. As a result, some sources prefer the use of other terms such as Long-Reach Optical Access Network [2–4]. However, today the name LR-PON is most widely used and will therefore be used in this dissertation.

Throughout the years there were multiple demonstrations showing the vi-

ability of LR-PONs. In 1997 there was a first lab demonstration of the SuperPON, which had a 2048 split and a 100 km reach [5]. A more recent proof-of-concept was shown in 2007 [3], when a 10 Gbit/s, 100 km reach LR-PON was demonstrated, serving 1024 users.

3.1.1 Impact on the ONU

The main advantage of an LR-PON stems from the elimination of a considerable part of the traditionally required equipment, while the equipment that is left in the network is shared by a much larger number of users. This effectively reduces the cost per user of the optical network, but has a major impact on the performance requirements of the ONU. This is explained as follows.

Cost-wise, it is advantageous to serve as high a number of users as possible with a single OLT. The OLT is shared among all users, which results in a cost per user decrease. However, in such a scenario the bandwidth per user also diminishes: the available bandwidth must be divided over a larger user base. To cope with this undesirable consequence, an OLT that supports a higher line rate should be installed. This implies a cost and power penalty, but thanks to the high number of users, the per-user implications are negligible.

However, since the OLT is transmitting at a higher line rate, all ONUs should also be updated to support this higher line rate. This results in a cost increase, both financially and power-wise, that is not shared as was the case for the OLT. The cost per ONU therefore dramatically rises when trying to provide the same user bandwidths.

Furthermore, the desired increased reach translates to a larger required optical budget. While this is partly accomplished by inserting optical amplifiers or Optical-Electrical-Optical (O-E-O) repeaters in the network [6], typically ONUs with higher sensitivity are required to provide the higher optical budget, which typically translates in more expensive or power-hungry devices.

3.1.2 LR-PON Conclusion

The core idea of an LR-PON certainly shows big advantages: per-user cost reduction thanks to a higher split ratio and network simplification attributed to a longer reach are things all next generation networks would benefit from. However, the implications on the ONU should not be ignored. Ideally, a so-

lution would offer this higher split ratio and longer reach, without increasing the performance requirements on the ONU equipment.

3.2   Taking BiPON to the next level

While Chapter 2 showed the Bit-interleaving PON achieving impressive power consumption reductions, there is no reason to stop where BiPON did. Due to the increasing demand for higher bandwidths [7], complemented with the expected explosion of connected devices [8], power consumption of our communication networks will rise to unacceptable levels, even when applying currently developed energy saving techniques. Therefore, research to keep lowering the networks' power consumption is indispensable. Since BiPON demonstrates such remarkable savings, this novel transmission protocol serves as a good starting point to continue the work.

3.2.1   Power deficiency of BiPON

The power savings realized with BiPON were achieved by focusing on the ONU, and more specifically on the *electronics* of the ONU. It was a logical point of focus, as careful investigation of the XG-PON processing chain revealed a blatant power deficiency there. With the introduction of the BiPON protocol, this power deficiency has been tackled. The success of the BiPON solution was the result of a design methodology that was focused on *one* goal: reducing the process rate of the received data as early in the chain as possible.

When analyzing the solution offered by the demonstrated 10 Gbit/s BiPON [9, 10], we find that methods have been developed within BiPON to reach this goal for the electronics of the ONU. Unfortunately, these mechanisms have had no effect on the optical part of the ONU. Even though an ONU in a 10 Gbit/s BiPON supports a maximum user rate of 1.25 Gbit/s, the optical components on the ONU side are still required to support 10 Gbit/s. This is of course undesirable, as it goes without saying that a 10 Gbit/s optical receiver is more power-hungry than a 1.25 Gbit/s one. Moreover, the cost of a 1.25 Gbit/s optical receiver is significantly lower than its 10 Gbit/s counterpart. While this shows that it would be beneficial to be able to use 1.25 Gbit/s optical components instead of 10 Gbit/s, the BiPON solution does not provide any means to do so.

3.2.2 Lack of suitable optical components

To bring the advantages of BiPON from electronics to the optics, an optical receiver that can directly sub-sample the incoming data stream is needed. Unfortunately, since the field of photonics is relatively new, optical components today have not reached the same maturity and possibilities as electronics, where very complex systems can be built with a multitude of tools and decades of experience in the field. This means that, for now, we only have relatively simple optical processing components at our disposal.

More specifically, today, no sub-sampling optical receivers are available and it is therefore not possible to apply the same techniques used for the electronics in a BiPON receiver when tackling the optical components.

3.2.3 Leveraging the complexity of electronics

Since no optical solution for the power deficiency in BiPON is to be expected any time soon, other options should be explored. The approach taken in this dissertation is to leverage the mature field of electronics to help the field of photonics where it is still lacking today. The research presented in this dissertation describes such a solution, called the Cascaded Bit-interleaving PON (CBi-PON). CBi-PON is a multi-level BiPON implementation, with an associated CBi-PON Protocol which is based on the BiPON Protocol.

3.3 Cascaded Bit-Interleaving PON: Concept

While the concept of an LR-PON shows a lot of potential both in terms of cost reduction as in power reduction of the network, it has serious implications on the required performance of the ONUs in the network. As there is a vast number of ONUs present in the network, these implications strongly reflect on the network with respect to cost and power consumption.

Since the methodology used to conceive BiPON rendered such impressive results, it was considered a promising approach to come up with a network design that would maximally exploit the potential of the LR-PON. Applying the same methodology is achieved by approaching the LR-PON with focus on one goal: *reducing the processing rate of the received data as early in the chain as possible*.

Since LR-PONs make use of amplified splitters, the network architecture already takes into account the presence of powered nodes throughout the network. When trying to reduce the data rate as early in the chain as

possible, these active nodes present an interesting opportunity: instead of just splitting the optical signal, some processing can be added. Therefore, the amplified splitters in the LR-PON are replaced by Optical-Electrical-Optical (O-E-O) converters that perform processing in the electrical domain. This processing allows a lower output data rate, effectively reducing the processing rate of the received data very early in the chain, which was our main goal. Furthermore, every added O-E-O converter allows to further reduce the data rate. It has to be noted that the cost and power consumption of these converters is shared by all end-users and therefore has limited impact on the per-user performance of the network.

This reasoning is what ultimately led to CBi-PON, which is essentially an LR-PON consisting of multiple bit-interleaved levels. The CBi Protocol, a derivative of the BiPON Protocol, was developed to provide support for multiple levels and to enable a flexible dynamic bandwidth allocation across the different levels. As it turns out, by combining the concept of the LR-PON with BiPON, we can eliminate most of the concerns that rise when assessing the original LR-PON concept.

The remainder of this chapter will introduce the reader to the network infrastructure of a CBi-PON and give an overview of the devices that are essential in such a network. Subsequently, each of these devices will be presented, discussing their function within the network and clarifying their operation. Finally, the CBi Protocol will be outlined. In conclusion, the advantages of a CBi-PON over other network topologies will be highlighted.

3.3.1   CBi Network Topology

On each level in a CBi-PON, CBi End-ONTs act as leaf nodes of the Bit-interleaved PON on that level. This allows different types of external networks to connect to the level that is best suited for that network's requirements.

The CBi-PON itself is connected to the core mesh network by means of a Core Router. The CBi Interleaver (CBi-I) takes care of the protocol translations necessary to let the CBi-PON interface with the core mesh network.

3.3.2   Downstream rates

To simplify the prototype design, we introduced a fixed relation between the downstream rates of adjacent CBi Levels: the downstream rate of a CBi Level is defined as a quarter of that of the higher CBi Level, as is
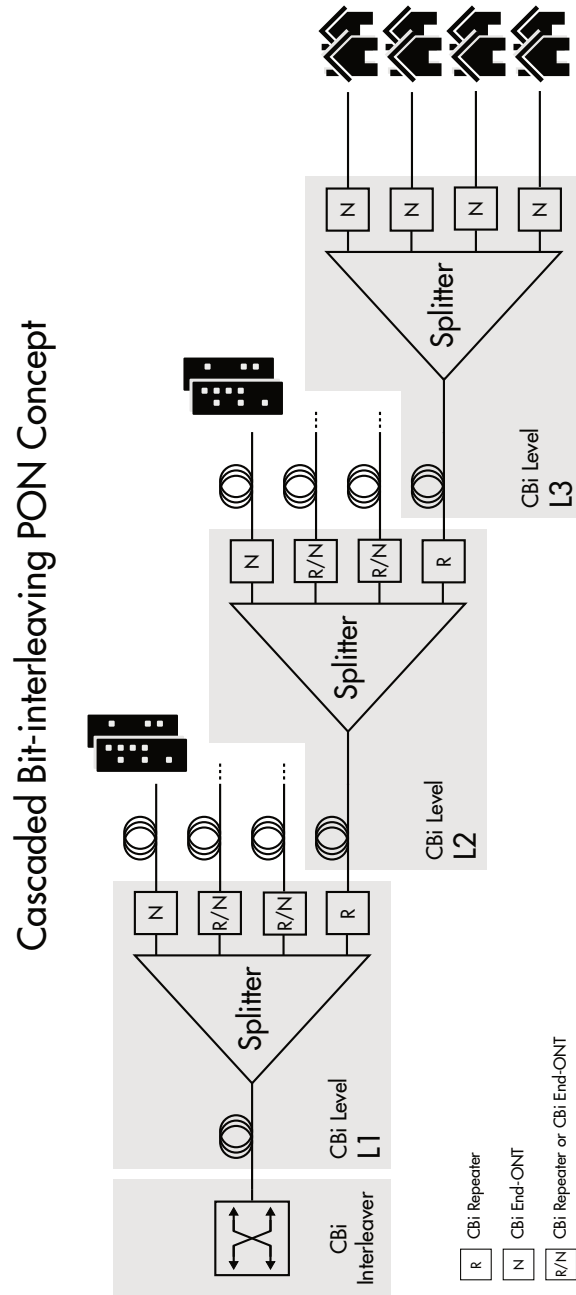
Figure 3.2: Cascaded Bit-Interleaving PON Network Architecture

symbolically expressed in Table 3.1. In principle, the CBi-PON concept supports other ratios between the downstream rates in adjacent CBi Levels as well.

	L_x **BiPON**	L_{x+1} **BiPON**
Rate	R_x	$R_{x+1} = \frac{1}{4}R_x$

Table 3.1: Downstream data rate relations in CBi PON

This rule determines the line rates in each CBi Level once the highest level downstream rate has been selected. Typically, the downstream rate at the highest level will be chosen high enough to provide the end users with a sensible bandwidth despite the high split ratio of the network.

3.3.3 CBi Devices

This section elaborates on the operation of the different CBi Devices essential to a CBi-PON. There are 3 CBi-PON specific devices required to build a CBi-PON: the CBi Interleaver, the CBi Repeater and the CBi End-ONT.

(a) CBi Interleaver Operation

(b) CBi Repeater Operation

(c) CBi End-ONT Operation
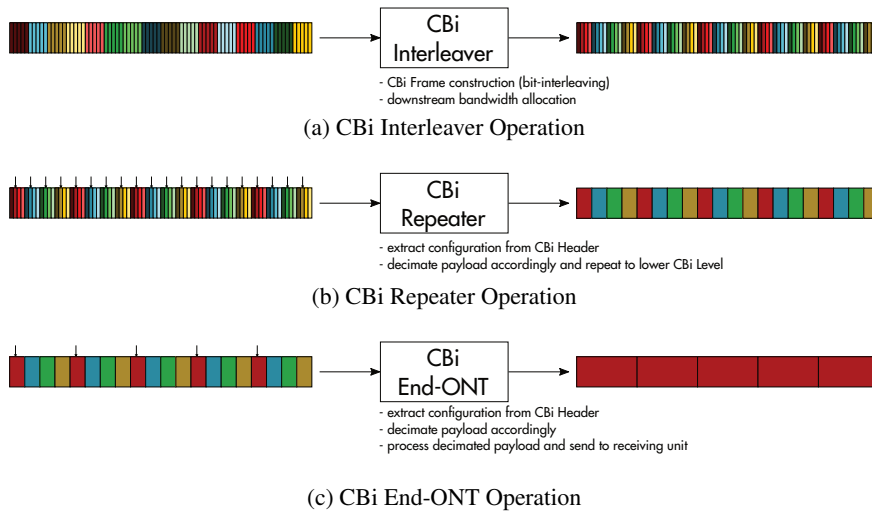
Figure 3.3: CBi Devices Operation

3.3.3.1   CBi Interleaver

The CBi Interleaver is the gateway between a CBi-PON and any other network: it takes on the role of the OLT in a typical LR-PON. While traffic in a CBi-PON is sent bit-interleaved, this is not the case in the other networks, which is why a translator device is required. Furthermore, the specific structure of the CBi Frame, as will be discussed in Section 3.3.4, requires the CBi Interleaver to correctly construct the frame header. Finally, the CBi Interleaver needs to take on the downstream bandwidth scheduling, as bandwidth allowance is communicated through the CBi Frame Header.

The CBi Interleaver has two interfaces. On the network side, typically where the core network resides, there is a port running a standard Network-to-network interface (NNI), while at the line side, the side of the CBi-PON, there is a port running the custom CBi Protocol.

The operation of the CBi Interleaver is depicted in Figure 3.3a. In the downstream direction, the CBi Interleaver receives typical packet-based traffic and does the required processing to send the data into the CBi-PON: a CBi Frame is constructed, naturally in a bit-interleaved fashion, and all data is arranged according to their destination in the CBi-PON.

Like in BiPON, the CBi Protocol is only applied on the downstream traffic, because the techniques that accomplish the desired energy savings in the downstream are pointless in the upstream direction.

3.3.3.2   CBi Repeater

The CBi Repeater (CBi-R) is the CBi Device that connects adjacent CBi Levels. Operating in the downstream direction, the CBi Repeater receives CBi Frames from the higher CBi Level. Subsequently, the CBi-R extracts the repeater configuration from the CBi Frame Header based on the CBi Repeater identifier: Repeater/End-ONT Identifier (RNID). This configuration allows the CBi Repeater to select the appropriate part of the payload and to repeat it to the lower CBi Level.

It is important to note that the CBi-R operation is deliberately kept as simple as possible: no descrambling or other complicated processing steps are required. Keeping the operation as close to a pure decimation step as possible allows to reduce the power consumption of the device.

3.3.3.3   CBi End-ONT

The CBi End-ONT acts as a leaf node of a CBi Level in the CBi-PON: it terminates the CBi-PON on that level and provides an interface to the user network. Note that a CBi End-ONT can be used in any of the CBi Levels and is not restricted to the lowest level.

Compared to the CBi Repeater, the CBi End-ONT has a much higher processing complexity: upon reception of a CBi Frame, the header is parsed to correctly configure the receiver which enables appropriate processing of the payload and extraction of the desired data. This is a complex process, which is described in more detail in Section 4.

3.3.4   CBi Frame Composition

Data transmitted over a CBi-PON is encapsulated in CBi Frames, based on the BiPON Frames introduced in Chapter 2. The largest part of a CBi Frame is the payload section, which is the actual data being sent. However, the CBi Protocol adds some overhead since the payload section is preceded by the header section. The CBi Header contains the necessary information to correctly receive, process and interpret the incoming CBi Frame.

The CBi Header is in fact a bit-interleaved compound of H *header lanes*, with H dependent on the downstream rate of the CBi Level the frame corresponds to. For example, a CBi Frame used in a 40 Gbit/s CBi Level will have 1024 header lanes, while a CBi Frame used in a 2.5 Gbit/s CBi Level will have only 64 header lanes. The number of header lanes determines the maximum split ratio, as there should be a disctinct header lane for each receiver.

A CBi Device will therefore initiate reception of a CBi Frame with an aggressive decimation of the bits with a decimation factor equal to the number of receivers supported in the CBi Level. As a result, one of the header lanes is recovered, allowing to extract the configuration information from the CBi Header. Of course, the decimation factor is re-adjusted once the full configuration of the CBi Device is completed, to receive the payload at the desired rate.

Every Header Lane consists of 3 distinct fields: SYNC, RNID, and BWMAP. The synchronization (SYNC) field is required to let the CBi Device synchronize with the incoming CBi Frame. The Repeater/End-ONT Identifier (RNID) is used to address the CBi Device. Finally, the Bandwidth Map
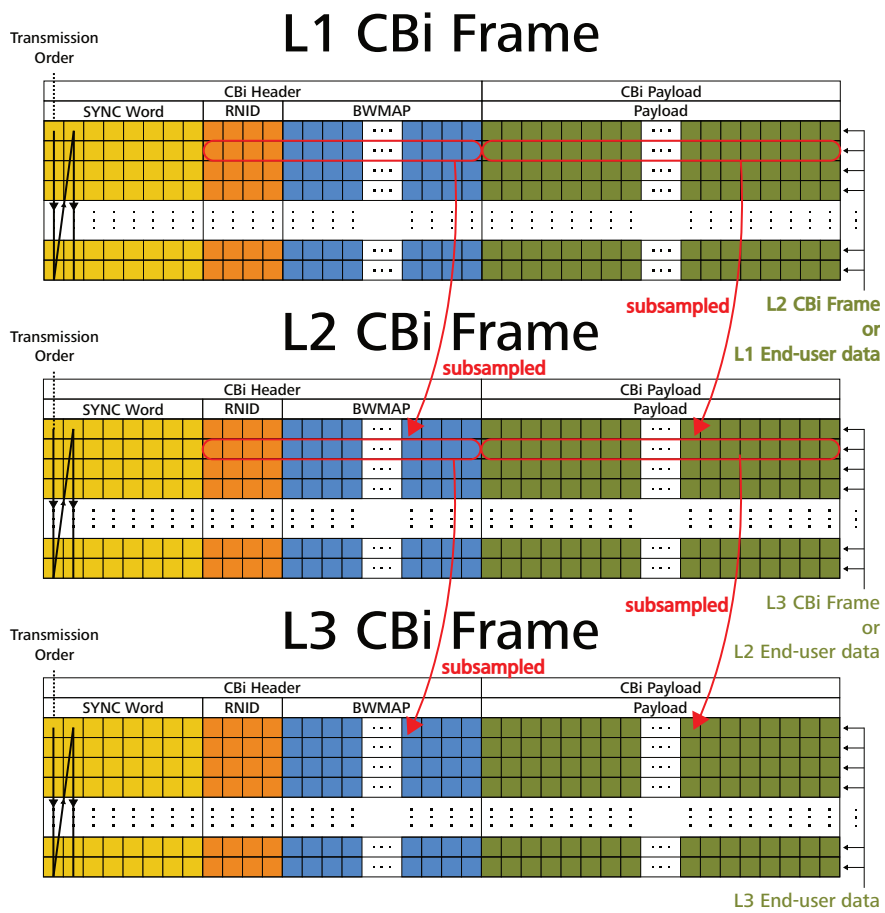
Figure 3.4: CBi Frame Composition

(BWMAP) contains the bandwidth allocations for each CBi Device, and provides the possibility to distribute upstream bandwidth allocations.

While the CBi Frames are based on the BiPON Frames, there is a very important difference. CBi Frames for any CBi Level that is not the lowest in the CBi-PON consist of interleaved lower-level CBi Frames. The header of the lower-level CBi Frame is retrieved by sub-sampling the header, while the payload is extracted by sub-sampling the payload. Note that the information for the CBi Repeater is contained in first header lane, which is reserved for this, and as a result all other header lanes only contain header information for End-ONTs. In other words, when an L_x CBi Repeater receives a CBi Frame, the payload that is extracted is in fact the payload of an L_{x+1} CBi Frame, to be further processed by an L_{x+1} CBi Repeater or End-ONT. This payload is then preceded by a header created by a sub-sampling of the incoming header. A visualization of this cascaded encapsulation is given in Figure 3.4.

Although all fields in the CBi Frame are the same for all CBi Levels, there are differences in the physical length of the fields depending on the CBi Level and the nominal rate. Take for example a CBi-PON consisting of three levels with the highest CBi Level running at at line rate of 40 Gbit/s. The field lengths for such a CBi-PON are summarized in Table 3.2. As can be seen in the table, the Frame duration is fixed as $125\,\mu$s by design. The CBi Header has a fixed length in terms of bits, which means it is the number of bits in the payload that has to vary with the CBi Level.

3.3.4.1 CBi Header: SYNC Field

The SYNC field contains a SYNC word, which allows the receiver to synchronize with the incoming frame. The SYNC word is a 32-bit fixed pattern: either `0x1840 FD59` or `0xE7BF 02A6`. These two words are used in an alternating fashion lane-to-lane to ensure DC-balanced transmission, since they are each other's complement, as shown in Table 3.3.

The initial header alignment is achieved by searching for the SYNC word on the downsampled bitstream. Once the bitstream has been recognized, alignment is achieved and further processing can commence.

3.3.4.2 CBi Header: RNID Field

Every receiver in a CBi-PON has a unique identifier. Since the identifier can be used to identify both a CBi Repeater or a CBi End-ONT, it is called the Repeater/End-ONT Identifier (RNID).

	L1	L2	L3
Rate (Gbps)	40	10	2.5
Header Lanes	1024	256	64
SYNC length (octets)	4	4	4
RNID length (octets)	2	2	2
BWMAP length (bits)	76	76	76
Payload length (octets)	606208	151552	37888
Frame length (octets)	622080	155520	38880
Header duration (μs)	3.18930	3.18930	3.18930
Payload duration (μs)	121.81070	121.81070	121.81070
Frame duration (μs)	125.00000	125.00000	125.00000

Table 3.2: Line rate relations in CBi PON

	SYNC Word
Even (hex)	0xE7BF 02A6
Odd (hex)	0x1840 FD59
Even (binary)	0b1110011110111111000000101010110
Odd (binary)	0b0001100001000000111111010101011001

Table 3.3: SYNC Words

The RNID in a particular header lane states the addressed receiver. Furthermore, if the receiving CBi Device has an RNID that is different from the one received, the required decimation offset can be calculated to correctly receive the data. Once the RNID in the received header lane matches the RNID of the receiving CBi Device, the processing of the CBi Frame can continue.

The RNID is a 2-octet field (16 bits), of which the 2 most significant bits are reserved. The 10 least significant bits are the LaneID, which for odd-parity lanes is encoded as a negative value represented in 2's complement format. Finally, the remaining 4 bits are used to ensure DC-balancing. The encoding of the RNID is summarized in Table 3.4.

	Reserved	DC-balancing	LaneID
Even-parity lane	00	0000	xxxxxxxxxx
Odd-parity lane	11	1111	not(xxxxxxxxxx)+1

Table 3.4: RNID Encoding

3.3.4.3   CBi Header: BWMAP Field

The BWMAP field consists of 76 bits of which the first 4 bits form a Flag subfield to indicate the presence of each of the consequent subfields: Downstream Bandwidth Map (DSBWMAP), Upstream Bandwidth Map (USBWMAP) and OAM.

Bit	Name	Description
1	DS_flag	1 = Presence of DSBWMAP
2	US_flag	1 = Presence of USBWMAP
3	OAM_flag	1 = Presence of OAM
4	undefined	Reserved

Table 3.5: BWMAP Flag Subfield

DSBWMAP Subfield

The DSBWMAP, 16 bits long, contains a doublet (R_k, O_k) for every receiver k that determines how to sample the payload section. R_k corresponds to the decimation rate and O_k to the offset with which the payload should be sampled. R_k consists of 4 bits indicating the down-sampling rate as shown in Table 3.6. The other 12 bits are assigned to the offset value O_k, which is used together with the RNID to determine the initial state of the descrambler and the location of the first payload bit in the CBi Frame.

$\mathbf{R}_k$	Downsampling Rate
00xx	1/4
01xx	1/8
10xx	1/16
11xx	1/32

Table 3.6: DSBWMAP Downsampling Rates

USBWMAP Subfield

The USBWMAP subfield, 16 bits long, has been included in the CBi Protocol to allow upstream bandwidth allocation. The first 8 bits can be used to indicate a time slot, while the last 8 bits indicate the duration of the upstream burst.

OAM Subfield

The Operations, Administration and Management (OAM) subfield does not have a fixed length, which means it could take up any number of the remaining 40 bits of the BWMAP field, and it starts immediately after the DSBWMAP or USBWMAP, if present.

An example of an OAM command is the *sleep* command, which tells a CBi Device to enter the sleep state for a number of frames, indicated by the data bits sent in the OAM subfield.

3.3.4.4   Scrambling

To ensure sufficient transitions in the transmitted data, the content of the BWMAP field and of the payload section of each CBi Frame is scrambled using a frame-synchronous additive scrambling polynomial: $1 + x^{-18} + x^{-23}$. This pattern is added modulo to the downstream data (i.e. XORed). The scrambling state is reset to a fixed pattern of all 1's at the first bit

following the SYNC word, to assure a deterministic scrambling state that allows descrambling upon reception with a simple XOR operation.

3.4 Cascaded Bit-Interleaving PON: 3-Level Instantiation

While the previous section presented the concept of a Cascaded Bit-interleaving PON, this section discusses an instantiation of the concept that is suitable to serve as a next generation network and is possible to build with currently available equipment.

3.4.1 3-Level Cascaded Bit-interleaving PON

The CBi concept is scalable to any number of levels, yet the research presented here is focused on a 3-level implementation, since this is the most in line with today's availability of devices and end user requirements.

Figure 3.5 shows the 3-level instantiation of a CBi-PON: L1, L2 and L3 are self-contained Bit-interleaved PONs. In agreement with the CBi concept, these levels are connected to each other by CBi Repeaters, and each level is possibly terminated by CBi End-ONTs acting as leaf nodes.

3.4.1.1 Rates

CBi Level	L1	L2	L3
Downstream Rate	40 Gbit/s	10 Gbit/s	2.5 Gbit/s

Table 3.7: Rates in proposed CBi-PON

In the 3-level proof-of-concept presented here, the L1 downstream line rate is chosen to be 40 Gbit/s to ensure sensible bandwidths while supporting a high split ratio. Following the rules regarding the rates within the CBi-PON results in the rates shown in Table 3.7.

3.4.2 CBi Frame Configuration

While the generic structure is the same, each CBi Level corresponds to a CBi Frame with a slightly different configuration. For each CBi Level, the implementation specifics are annotated on Figure 3.6.
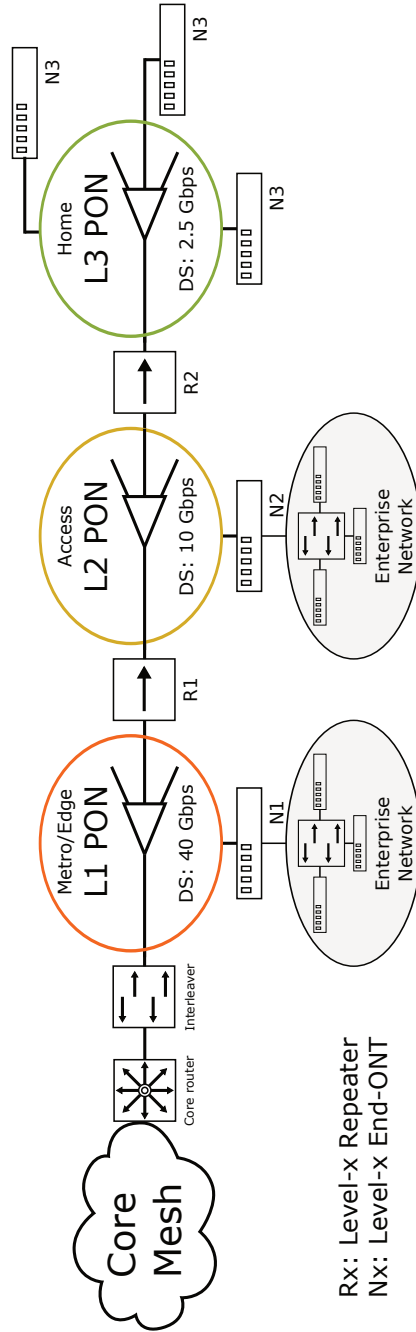
Figure 3.5: 3-Level Cascaded Bit-interleaving PON Instantiation
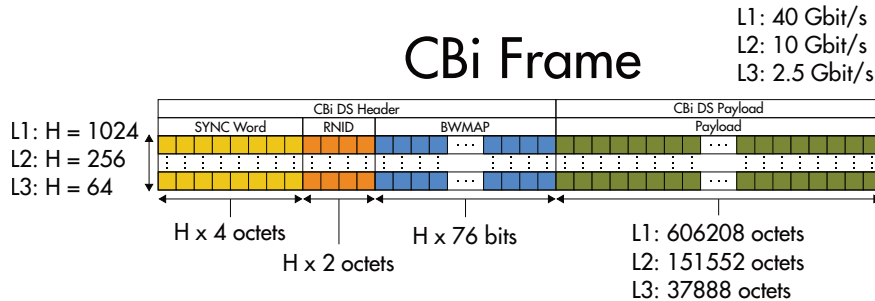
Figure 3.6: CBi Frame Configuration

3.4.3   CBi Interleaver Implementation

A block diagram of a proof-of-concept CBi Interleaver is shown in Figure 3.7. In the proof-of-concept network, focus is on the CBi Repeaters and End-ONTs. Therefore, the CBi Interleaver is implemented using only commercial components.

An FPGA (e.g. Altera DEV-5SGX) is used to generate the necessary traffic, providing $4\times$ 10 Gbit/s. These 4 streams are then multiplexed (e.g. by a HMC847LC5) to form a 40 Gbit/s signal that is sent to the Laser Driver (LD) which drives the Transmitter Optical Sub-Assembly (TOSA).
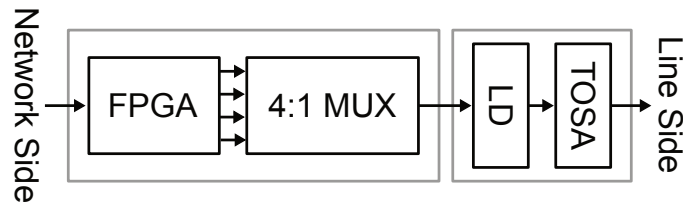
Figure 3.7: CBi Interleaver Block Diagram

3.4.4   CABINET Application-Specific Integrated Circuit

To show the potential power savings possible in the network, an Application-Specific Integrated Circuit (ASIC) is required to implement a generic CBi Device. Implementation would be possible on an FPGA, however, this would require commercial Clock-and-Data Recovery (CDR) circuits and would increase the power consumption. Furthermore, due to the vast amount

of CBi Repeaters and CBi End-ONTs in a CBi-PON, it makes sense economically to create a dedicated chip which is less expensive than FPGAs when producing in high volumes.

Therefore, in this research the CAscaded Bit-Interleaving eNd tErmination/repeaTer (CABINET), a generic CBi Device ASIC, was designed and fabricated. More details on the design of the CABINET can be found in Chapter 4.

3.4.5 CBi Repeater Implementation
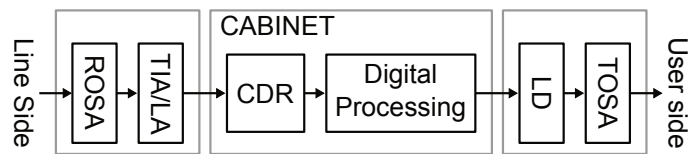
Figure 3.8: CBi Repeater Implementation Block Diagram

The CBi Repeater (Figure 3.8) makes use of a commercial Receiver Optical Sub-Assembly (ROSA) to receive the incoming data, which is followed by a commercial Transimpedance Amplifier (TIA) and Limiting Amplifier (LA) to convert the optical signal to the electrical domain. The CABINET ASIC then recovers the clock and the data and performs the limited processing needed in a CBi Repeater: extracting the header information and repeating the decimated payload. The output stream of the CABINET is then sent to the LD which drives the TOSA, which are again off-the-shelf components.

3.4.6 CBi End-ONT Implementation
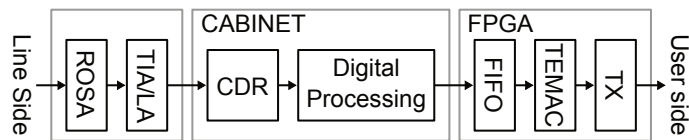
Figure 3.9: CBi End-ONT Implementation Block Diagram

The input stage of a CBi End-ONT (Figure 3.9) is the same as a CBi Repeater: a ROSA and a TIA/LA combination. However, the CABINET is

configured as a CBi End-ONT, and does more elaborated processing of the CBi Frame, resulting in a descrambled payload that is sent to the FPGA. The FPGA is used to provide a standard interface to the network the CBi End-ONT connects to.

3.5   Conclusion

In this chapter the concerns for a Long-Reach PON were expressed, as well as the issues that arise when trying to scale the BiPON solution. A solution to these concerns was proposed in the form of the Cascaded Bit-interleaving PON concept. The CBi-PON network topology, CBi Devices and the CBi Frame were introduced, and finally a 3-level proof-of-concept CBi-PON was presented. The implementation of the different CBi Devices was discussed and showed the need for a generic CBi Device ASIC: the CABINET.

References

[1] D.B. Payne and R.P. Davey. The future of fibre access systems? *BT Technology Journal*, 20(4):104–114, 2002.

[2] RP Davey, D Nesset, A Rafel, DB Payne, and A Hill. Designing long reach optical access networks. *BT Technology Journal*, 24(2):13–19, 2006.

[3] D.P. Shea and J.E. Mitchell. A 10-Gb/s 1024-Way-Split 100-km Long-Reach Optical-Access Network. *Lightwave Technology, Journal of*, 25(3):685–693, March 2007.

[4] H. Song, B. W. Kim, and B. Mukherjee. Long-reach optical access networks: A survey of research challenges, demonstrations, and bandwidth assignment mechanisms. *IEEE Communications Surveys Tutorials*, 12(1):112–123, First 2010.

[5] Jan Vandewege, Xing-Zhi Qiu, Brecht Stubbe, Chris Coene, Peter Vaes, Wei Li, Jan Codenie, Claire Martin, H Slabbinck, Ingrid Van de Voorde, et al. A lab demonstration of a SuperPON optical access network. In *Broadband Communications*, pages 69–80. Springer, 1998.

[6] R.P. Davey, P. Healey, I. Hope, P. Watkinson, D.B. Payne, O. Marmur, J. Ruhmann, and Y. Zuiderveld. DWDM reach extension of a GPON to 135 km. In *Optical Fiber Communication Conference, 2005. Technical Digest. OFC/NFOEC*, volume 6, pages 3 pp. Vol. 5–, March 2005.

[7] Cisco Visual Networking Index. Forecast and Methodology, 2014-2019 White Paper. Technical report, Technical Report, Cisco, 2015.

[8] Cisco Consulting Services. The Internet of Everything (IoE) - Connections Counter, 2013.

[9] Christophe Van Praet, Hungkei Chow, Dusan Suvakovic, Doutje Van Veen, Arnaud Dupas, Roger Boislaigue, Robert Farah, Man Fai

Lau, Joseph Galaro, Gin Qua, et al. Demonstration of low-power bit-interleaving TDM PON. *Optics express*, 20(26):B7–B14, 2012.

[10] Christophe Van Praet, Guy Torfs, Zuyi Li, Xin Yin, Dusan Suvakovic, Hungkei Chow, Xing-Zhi Qiu, and Peter Vetter. 10 Gbit/s bit interleaving CDR for low-power PON. *Electronics letters*, 48(21):1361–1363, 2012.

4

Design of the CABINET ASIC

In the previous chapter, the concept of the Cascaded Bit-interleaving PON was introduced. It was concluded with the presentation of a 3-level implementation of the concept, which made use of a generic CBi Device: the CABINET ASIC. The CABINET was designed to serve both as a CBi Repeater and as a CBi End-ONT, while supporting operation at 40 Gbit/s, 10 Gbit/s and 2.5 Gbit/s.

This chapter discusses the design of the CABINET ASIC, starting with the system architecture and followed by an in-depth discussion covering the most critical building blocks of the Analog-Front End. Subsequently, the Medium Access Control (MAC) preprocessor operation is briefly clarified and finally the implementation of the Analog Back-End is presented.

4.1 System Architecture

The multi-mode, multi-rate nature of the CABINET ASIC makes it a fairly complex system. To cope with the design of such a complex system, it is subdivided in multiple processing blocks.

This subdivision results in the system architecture of the CABINET device, shown in Figure 4.1. It consists of three major sections, each having a clearly defined objective: (1) the Analog Front-End, (2) the MAC Preprocessor Core and (3) the Analog Back-End.

(a) Repeater Mode Data path

(b) End-ONT Mode Data path
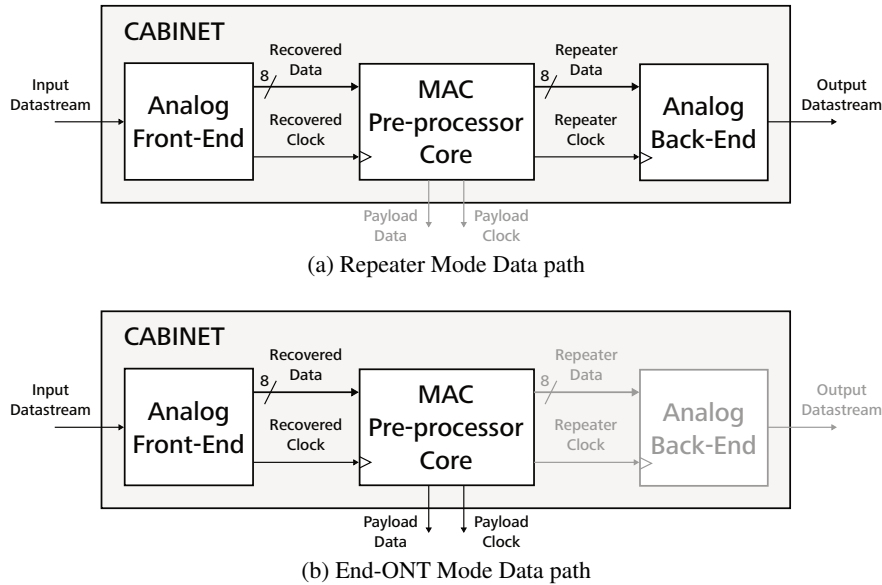
Figure 4.1: CABINET System Architecture

The objective of each of these processing blocks can be defined as:

Analog Front-End
Recovers the data and the clock from the incoming data stream.

MAC Preprocessor Core
Extracts the data to be repeated or the payload to be delivered.

Analog Back-End
Composes the output data stream from the repeater data.

4.1.1 Inputs

Regardless of mode or rate, the input of the CABINET is a data stream coming from an optical front-end. The optical front-end effectively performs an Optical-to-Electrical (O/E) conversion, of which the resulting data stream is used as the input for the CABINET device.

Even though the input will always come from an optical front-end, it is important to realize that the optimal optical front-end is highly dependent on the rate it should support. For example, the optical front-ends used in a CBi Level L1 (40 Gbit/s) requires high performance, expensive components, while for a CBi Level L3 (2.5 Gbit/s) a lower performance, and thus

less expensive, optical front-end is sufficient. Furthermore, noise and dispersion significantly increase with higher rates, which limits the maximum fiber length. Enabling a lower rate in the CBi Level L3 therefore allows the use of longer fibers and hence supports longer reach.

4.1.2   Outputs

While the CABINET has only one input port, there are two possible output ports, on account of the CABINET being a multi-mode device. The active output is determined by the configuration of the CABINET to act either as a CBi Repeater or as a CBi End-ONT, which is visualized in Figure 4.1.

Repeater Mode

When configured as a CBi Repeater, the MAC preprocessor extracts the part of the payload to be repeated based on the information from the CBi Header. This data is subsequently sent to the Analog Back-End, which constructs the output data stream to be transmitted by an optical module into the lower CBi Level, as depicted in Figure 4.1a.

End-ONT Mode

On the other hand, when the CABINET is configured to function as a CBi End-ONT (Figure 4.1b), the MAC preprocessor does a more extensive processing before delivering the payload and its associated clock to a Field Programmable Gate Array (FPGA). This FPGA is used to easily provide a standard Ethernet interface to the connected Local Area Network (LAN).

4.2   Analog Front End: Clock and Data Recovery

The Analog Front-End processing block of the CABINET device is responsible for recovering both clock and data from the incoming data stream. Subsequently, the data should be sent to the MAC Preprocessor Core to be processed. A schematic overview of the Analog Front-End is displayed in Figure 4.2. It is clear the Analog Front End consists mainly of a Clock-and-Data Recovery (CDR) block, which is preceded by an input buffer.
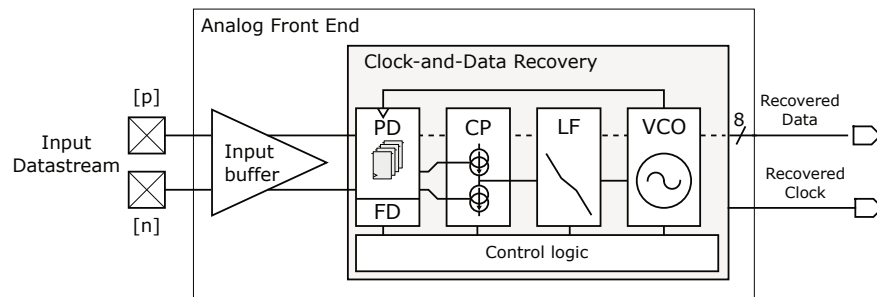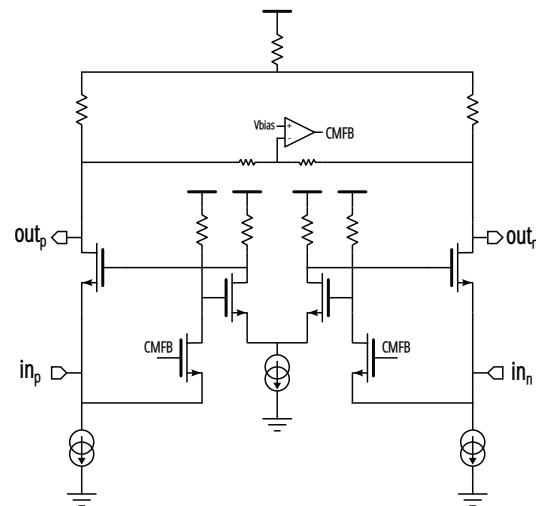
Figure 4.2: Analog Front-End

Figure 4.3: Input Buffer schematic

4.2.1   Input Buffer

The CABINET uses an input buffer (Figure 4.3) to provide the desired interface regarding voltage levels and impedance levels between the external signals and the internal circuitry. Furthermore, the input buffer can present a lighter load to the driving chip, which is particularly interesting in the case of high speed electronics such as the CABINET ASIC.
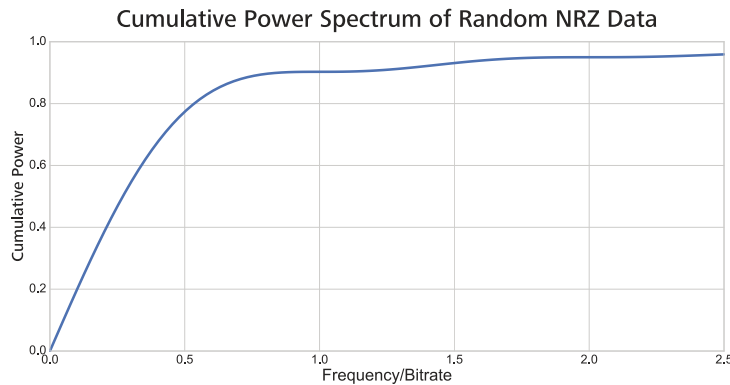
Figure 4.4: Cumulative power spectrum of random NRZ data [1]

The need of the CABINET to support input data streams up to $40\,\text{Gbit/s}$ determines the bandwidth requirement of the input buffer. Assuming an Non-Return to Zero (NRZ) modulation format, the bandwidth of the input buffer should be at least $30\,\text{GHz}$. This stems from observing Figure 4.4: for NRZ modulation, a bandwidth of $0.75\times$ the bit rate results in having received 93.6% of the total signal power. The high sensitivity of the sampler stage used in the CDR means a low voltage swing suffices to recover the input data, which significantly relaxes the gain requirement of the input buffer. Moreover, the optical module has already amplified the incoming signal to sensible voltage levels, further simplifying the design of the input buffer. The input buffer was implemented as a TransImpedance Amplifier (TIA) stage, since it is known to have a well-specified input impedance over a large bandwidth and can therefore be well matched to a $50\,\Omega$ source.

4.2.2   CDR Topology

The goal of the Analog Front-End is to prepare the input signal such that it can be processed by the MAC preprocessor. The heart of the Analog Front-End is the Clock-and-data Recovery (CDR) block.

This preparation consists of two main tasks:

1. Clock and data recovery

2. Rate reduction by means of deserialization

The CDR block used in the CABINET device combines both tasks by deserializing the data early on in the clock-and-data recovery loop. This combination has the additional advantage that some of the CDR blocks can operate at lower speeds, which is beneficial for power consumption.

4.2.2.1   Line rate dependence

While the principle of operation of the CABINET is not rate-dependent, the specific implementation of the CDR is. Consequently, several building blocks have multiple versions due to the multi-rate support of the CABINET. During the discussion of the building blocks, these differences will be highlighted when present.

4.2.2.2   Need for clock and data recovery

As the name suggests, Clock and Data Recovery performs two important functions: *clock recovery* and *data recovery*.

When transmitting NRZ modulated data over a channel, typically no clock is sent along with it. Sending a clock together with the data over the same channel would severely lower the spectral efficiency, while adding an extra channel just to send the clock would be overly expensive.

However, as we are dealing with digital communications, the receiver needs a clock to correctly sample the incoming data. This sampling of the incoming data is what is regarded as the *data recovery* and is dependent on the availability of a suitable sampling clock. While data recovery might seem trivial, the retiming function it performs is a crucial step, as it largely removes the jitter that was accumulated during transmission of the data stream.

This shows that to perform the crucial data recovery, an accurate sampling clock is indispensable. Since such a clock is not readily available at the receiver, we have to rely on clock recovery techniques to extract the clock that allows us to correctly sample the incoming data stream.

4.2.2.3 Need for deserialization

The need for deserialization is largely dictated by the choice of process technology, which in the case of the CABINET is a 40 nm Complementary Metal-Oxide-Semiconductor (CMOS) technology. This choice is justified by the large amount of digital processing, requiring a relatively large, and thus costly, chip area. The 40 nm node offers a good compromise between price and chip area. Furthermore, 40 nm CMOS offers devices that are sufficiently fast to accommodate the required high speed analog circuitry.

The MAC preprocessor is implemented by use of the Standard Cell Digital library provided by the foundry. This library is optimized to provide sufficient speed for typical applications (rarely above a few GHz), while minimizing area and power consumption. However, when considering an L1 CBi Repeater or L1 CBi End-ONT, the data stream that should be processed by the MAC preprocessor is clocked at 10 Gbit/s, which largely exceeds the operating capability of the available standard cells.

Due to this speed limitation, the MAC preprocessor has been modified to allow parallel processing. This adjustment allows the core to run at 1.25 GHz instead of 10 GHz, but requires parallel inputs, a requirement that is met by the deserialization operation in the CDR.

Note that this rate reduction is not equivalent to the rate reduction pursued within the spirit of the CBi-PON: it is merely a serial-to-parallel conversion to cope with the limited performance of the standard cells and does not incorporate a power consumption reduction. This is evident from the simplified equation of the dynamic power consumption of a digital cell (Equation 4.1): while the operating frequency is divided by 8, the number of cells multiplies by 8, resulting in a status quo in terms of power consumption.

$$P_{dynamic} = \frac{\alpha \cdot f \cdot C_L \cdot V_{dd}{}^2}{2} = 8 \times \frac{\alpha \cdot \frac{f}{8} \cdot C_L \cdot V_{dd}{}^2}{2} \qquad (4.1)$$

4.2.2.4 CDR Topology Selection

Years of research has produced a multitude of techniques to perform the desired recovery for multi-gigabit NRZ communication and all of them rely on the presence of sufficient data transitions, which forms the basis of the clock recovery.

An overview of these techniques is given in [2]. Each CDR topology can be classified into one of the following three main categories:

Feedback phase tracking
These include Phase Locked Loop (PLL), Delay Locked Loop (DLL), Phase Interpolator (PI) and Injection Locked (IL) topologies.

Oversampling without feedback phase tracking
An oversampling topology collects an excess of data samples and selects the best one as the data sample.

Phase alignment without feedback phase tracking
Examples are gated oscillators and high-Q bandpass filter structures.

An oversampling architecture requires a multitude of clock phases to be generated at the full rate, needs multiple high speed samplers and additionally requires a phase selection algorithm. Therefore, oversampling architectures are among the most power hungry CDR topologies and is therefore not suitable for the low power solution that the CBi-PON is pursuing. However, due to their fast acquisition properties, they are often used in burst-mode receivers.

Furthermore, even though this dissertation focuses on the downstream path, all CBi Devices are expected to support both down- as upstream communication. ONUs typically reuse the clock that is extracted from the downstream data traffic for the upstream transmission. This reuse has a significant impact on the jitter rejection requirements of the clock extraction. Therefore, due to their lack of jitter rejection, phase alignment topologies without feedback phase tracking are not acceptable for use in ONUs.

It is evident that for our application, the desired CDR topology will be from the feedback phase tracking category, but this still leaves us with a lot of options: Phase-Locked Loop, Delay-Locked Loop, Phase Interpolator or Injection Locked.

Both the Phase Interpolator and Injection Locked architectures require multiple *full-rate* clock phases [2], resulting in high power consumption and are therefore ruled out. A Delay-Locked Loop is only applicable to source-synchronous systems [2], where transmitter and receiver use the same clock source. While a DLL by itself is not useful for our application, dual-loop DLL/PLL topologies do exist that support asynchronous systems. These architectures combine the advantages of DLLs and PLLs, providing fast acquisition while avoiding jitter peaking, but this comes at a price: the dual-loop nature significantly complicates the system analysis and raises stability concerns. Furthermore, the presence of two loops is likely to require two

off-chip compensation capacitors, due to the large capacitors which are typically not implemented on-chip because of area restrictions increasing the total system cost.

As a result, the CDR topology chosen for the CABINET is a *Phase-Locked Loop based structure*.

An analog PLL was chosen as the underlying structure of the CDR. Furthermore, an external reference clock is used to simplify the frequency locking, since in the absence of an external reference, the frequency locking loop must rely on the incoming data stream.

4.2.2.5   PLL-based CDR Operation

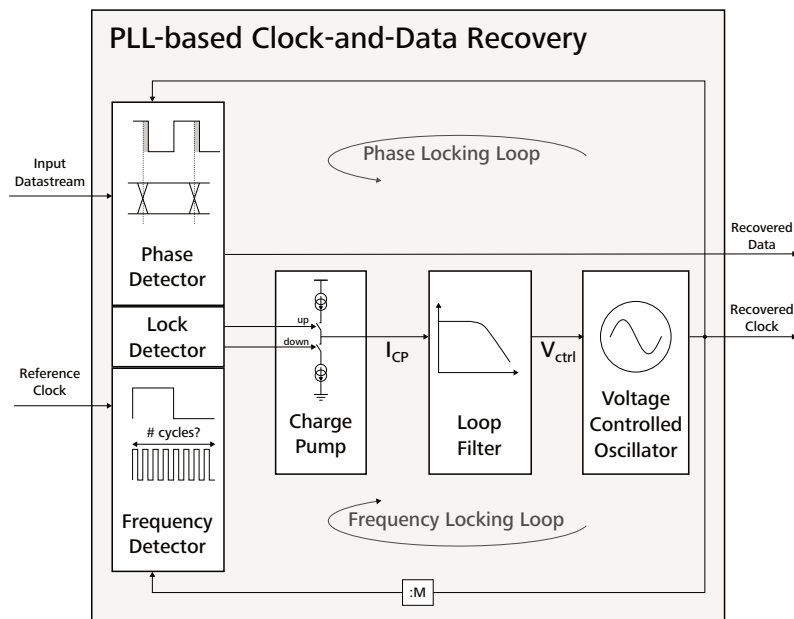Figure 4.5: Simplified block diagram of a PLL-based CDR using an external reference clock

In this section, the general principle of an Analog PLL-based CDR using an external reference clock is briefly explained. A simplified schematic is presented in Figure 4.5 and shows two loops: a Frequency Locked Loop (FLL) and a Phase Locked Loop (PLL).

Both loops are similar in operation and, consequently, share most of the

building blocks. The only difference is the use of either a Frequency Detector for the FLL and a Phase Detector for the PLL.

The goal of the Frequency Locked Loop is to adjust a Voltage Controlled Oscillator (VCO) such that the oscillation frequency is approximately equal to $M\times$ (e.g. $128\times$) the reference clock frequency. The reference clock frequency is chosen to be the line rate divided by M to match the incoming data stream.

In order to do so, the VCO output is divided by M and fed to the Frequency Detector, which determines if the VCO is either too fast or too slow, and sends the corresponding *up* or *down* signal to the Charge Pump (CP). The Charge Pump then accordingly either sources or sinks current into the Loop Filter, which gives rise to a change in control voltage at the VCO, causing the oscillation frequency to either increase or decrease accordingly.

Once the oscillation frequency is sufficiently close to the desired frequency, the FLL is said to be in lock. At this point, the FLL operation is halted and taken over by the Phase Locked Loop (PLL). This means that at any time, only one feedback loop is active.

The goal of the PLL is to continuously adjust the VCO frequency in order to phase-align the oscillator with the incoming data, such that the incoming data can reliably be sampled with the clock provided by the VCO.

The Phase Detector detects the phase difference between the incoming data and the oscillator output. Similar to the FLL, an *up* or *down* pulse is sent to the Charge Pump, although now the generation of this pulse is based on the phase difference. The Charge Pump output current is then transformed to a control voltage by the Loop Filter transfer function. This will either speed up or slow down the oscillator, resulting in the desired phase shift and ultimately aligning the oscillator clock with the incoming data.

4.2.2.6 A sub-sampling CDR

As has been noted throughout this dissertation, many of the advantages of the CBi-PON are attributed to the rate reduction early in the chain. Looking at the different CBi Levels, we remember from Chapter 3 that the rate in the lower CBi Level is chosen to be 1/4 of the upper CBi Level. This means only 1 out of 4 incoming bits is effectively repeated by a CBi Repeater. Likewise, a CBi End-ONT always uses a decimation factor of at least 1/4. Therefore, using the CDR to recover the complete incoming data stream would require us to discard 75% of the recovered bits, since every 3 out of 4 recovered bits is addressed to an other receiver. It is clear this is highly

inefficient. As a result, the CABINET is perfectly suited to use a true sub-sampling CDR: a CDR that recovers a clock at 1/4 of the incoming data, and recovers only 1/4 of the incoming data bits.

In the CABINET, this sub-sampling behavior was implemented for the 40 Gbit/s CDR. As this is the highest rate, the CABINET consumes the highest power in this configuration and therefore this is where the highest gain can be expected. The 10 Gbit/s and 2.5 Gbit/s CDRs have been implemented as regular, full-rate CDRs for the sake of simplicity and design time. The resulting CDR configuration for the CABINET is visualized in Figure 4.6.
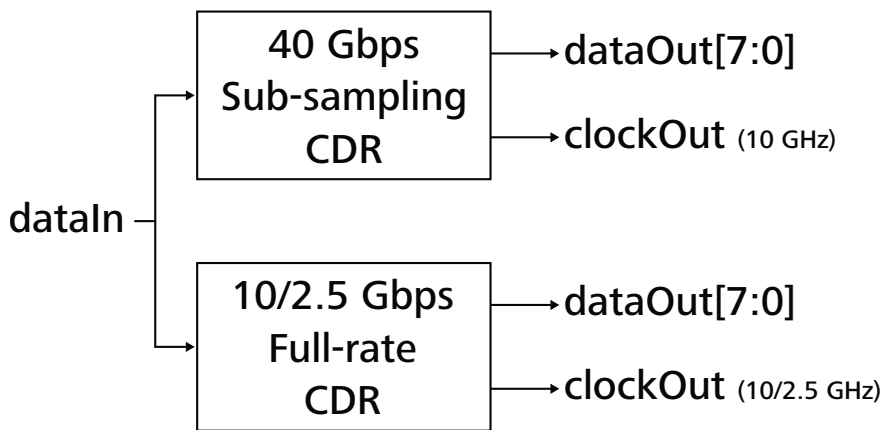

Figure 4.6: CABINET Multi-rate CDR Configuration

4.2.3 Phase Detector

The purpose of the phase detector is to compare the phase of the clock output of the VCO with the incoming data transitions. The phase detector transforms any phase misalignment in a corresponding error signal, which is then sent to the charge pump to adjust the phase alignment until the error signal is approximately zero.

Practically all PLL-based CDRs are subdivided in two categories based on the Phase Detector (PD) they use: (1) Linear CDRs making use of a *Linear* or *Hogge* PD [3], and (2) Bang-Bang CDRs making use of a *Bang-Bang* or *Alexander* PD [4]. The difference between the two types of PDs is visualized in Figure 4.7. A Bang-Bang PD outputs an error signal that only relates to the sign of the phase misalignment, while a Linear PD outputs an

error signal that is proportional to the phase error, and therefore relates to both the sign *and* the magnitude of the phase misalignment.
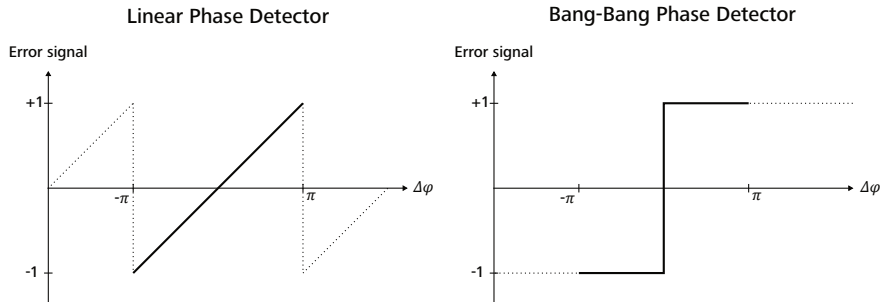

Figure 4.7: Linear vs Bang-Bang Phase Detector

The straightforward operation of the bang-bang phase detector generally leads to a clear-cut implementation and is therefore typically preferred [5] for high speed designs such as the CABINET. On the other hand, the extra information in the error signal of a Linear PD allows for less output jitter. However, the jitter requirement on the recovered clock in a CDR is not as stringent as for RF applications. For example, the GSM specification for Phase Noise is -130 dBc/Hz at 1 MHz offset from the carrier [6], while a CDR can typically sustain a Phase Noise as high as -86 dBc/Hz at 1 MHz for an oscillator running at 5 GHz (e.g. in [7]). As such, jitter is less of a concern in our application.

Moreover, the implementation of a Linear PD results in a Pulse-Width Modulation (PWM) signal at the output. At high rates, this makes it very challenging to obtain sufficient resolution. As a result, the benefits of the BB-PD outweigh those of the Linear PD and the CDR in the CABINET is implemented as a Bang-Bang CDR.

4.2.3.1 Phase Detector Architecture

The Bang-Bang Phase Detector used in the CABINET CDR consists of two main parts: the sampling stage and the Phase Detector Logic. These two stages implement the typical $2\times$ oversampling architecture of a BB-CDR. The Bang-Bang Phase Detector needs to identify the transitions in the incoming data stream and does so by sampling both at the rising and at the falling edge of the clock, effectively sampling at double the data rate. Based on 3 consecutive samples, the PD Logic can determine if the clock is either early or late with respect to the data. Subsequently, the output of

the PD Logic is used to drive the Charge Pump. This process is visualized in Figure 4.8.
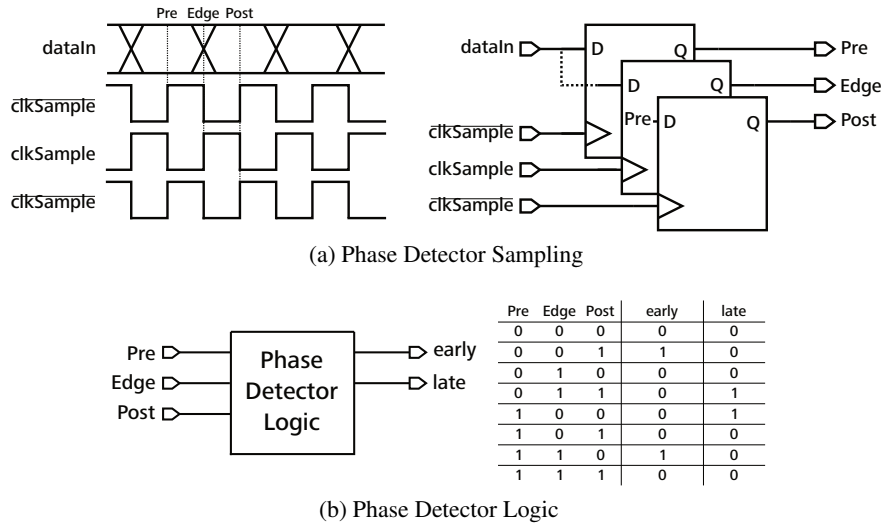

(a) Phase Detector Sampling

Pre	Edge	Post	early	late
0	0	0	0	0
0	0	1	1	0
0	1	0	0	0
0	1	1	0	1
1	0	0	0	1
1	0	1	0	0
1	1	0	1	0
1	1	1	0	0

(b) Phase Detector Logic

Figure 4.8: Bang-Bang Phase Detector Operation

	Sample Pre	**Sample Edge**	**Sample Post**
Channel 0	clkSample[7]	clkSample[0]	clkSample[1]
Channel 1	clkSample[1]	clkSample[2]	clkSample[3]
Channel 2	clkSample[3]	clkSample[4]	clkSample[5]
Channel 3	clkSample[5]	clkSample[6]	clkSample[7]

Table 4.1: Sampling clocks subset by Channel Selection

In case of the sub-sampling 40 Gbit/s CDR, only 1 out of 4 incoming bits is recovered, which means a 10 GHz clock suffices. However, to correctly identify transitions in the incoming 40 Gbit/s data stream, 8 equidistant clock phases are required, even though only 3 adjacent clock phases are used during operation. Depending on the selected channel, a different set of clock phases is used (Table 4.1). This modified scheme is summarized in Figure 4.9.
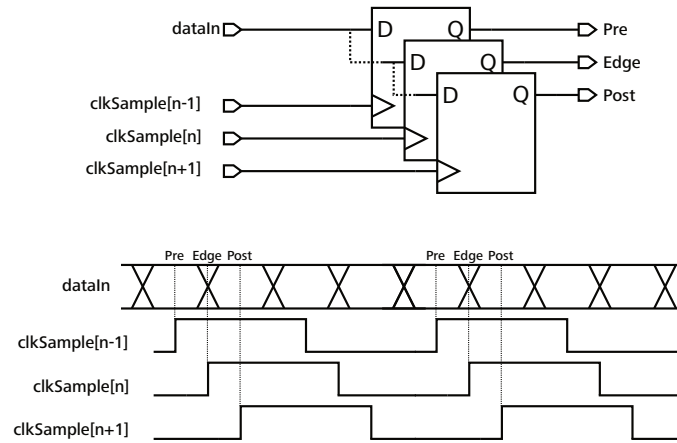
Figure 4.9: 40 Gbit/s Sub-sampling BB-PD Operation

4.2.3.2   Sampling Stage Implementation

The sampling stage consists of Deserializing Samplers: 8 for the 40 Gbit/s CDR (Figure 4.10) and 2 for the 10/2.5 Gbit/s CDRs (Figure 4.11). The Deserializing Sampler captures the input data stream and subsequently deserializes the high-speed single stream to 8 lower-speed streams. The sampling stage is implemented by means of a Sense Amplifier based Flip-Flop (SA-FF) as found in literature [8, 9] (Figure 4.12). A SA-FF has a fast sampling input followed by a slower regenerative section. This makes it an ideal choice for a sub-sampling stage, which needs to capture the high speed input data very quickly, but has relaxed requirements on the clock-to-output delay.

The SA-FF high speed sampler is followed by a tree of 1:2 deserializers that form the 1:8 deserializer, as depicted in Figure 4.13. To clock the deserializers, three clock dividers are used to generate the desired clock frequencies. The implementation of the 1:2 deserializer consists of two dynamic flip-flops clocked with the opposite clock edge, as presented in Figure 4.14.

4.2.3.3   Phase Detector Logic Implementation

The Phase Detector Logic is very straightforward and has been implemented as a digital block, which was synthesized and place-and-routed using the typical digital flow. Due to the deserializing operation of the
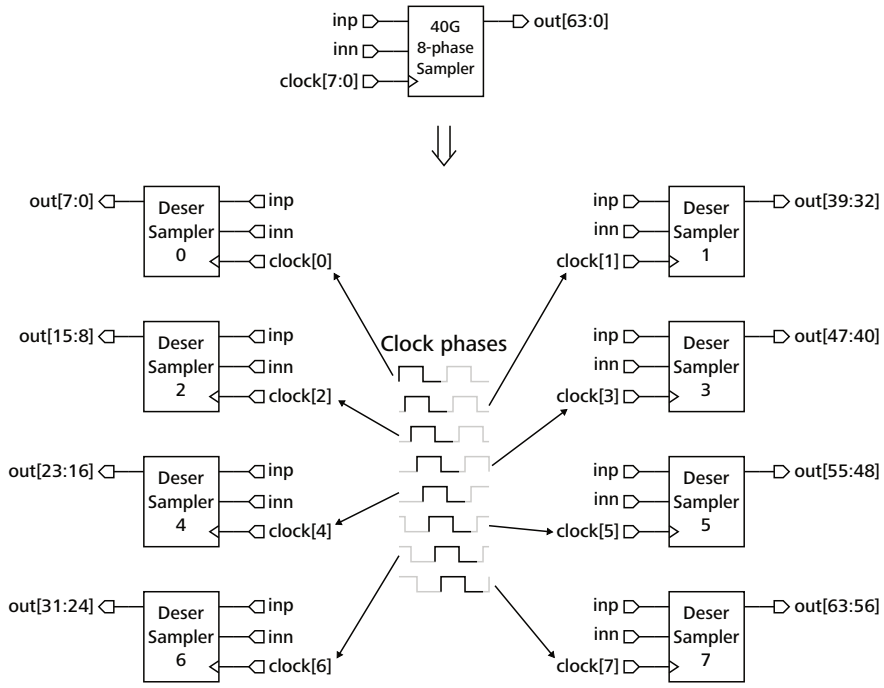
Figure 4.10: Samplers configuration for 40 Gbit/s CDR: 8 clock phases
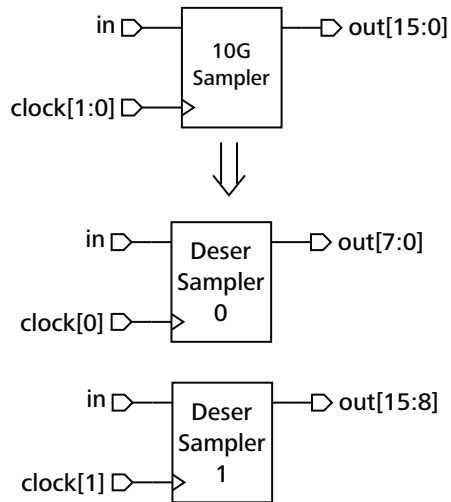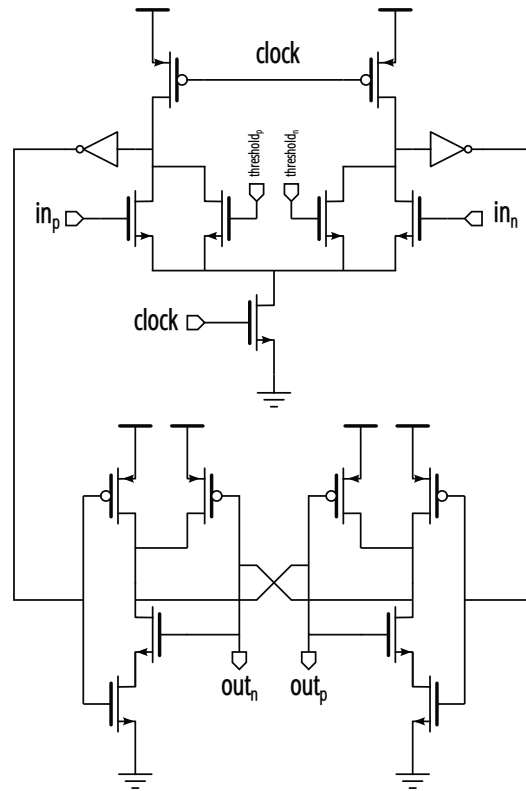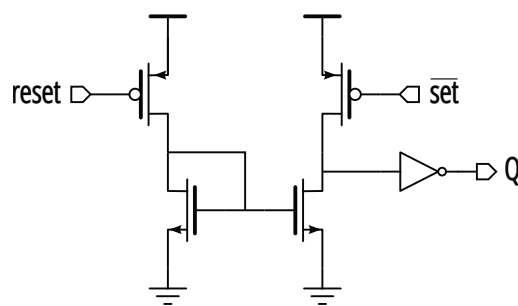
Figure 4.11: Samplers configuration for 10/2.5 Gbit/s CDR: 2 clock phases

(a) Sampling stage

(b) Latch stage

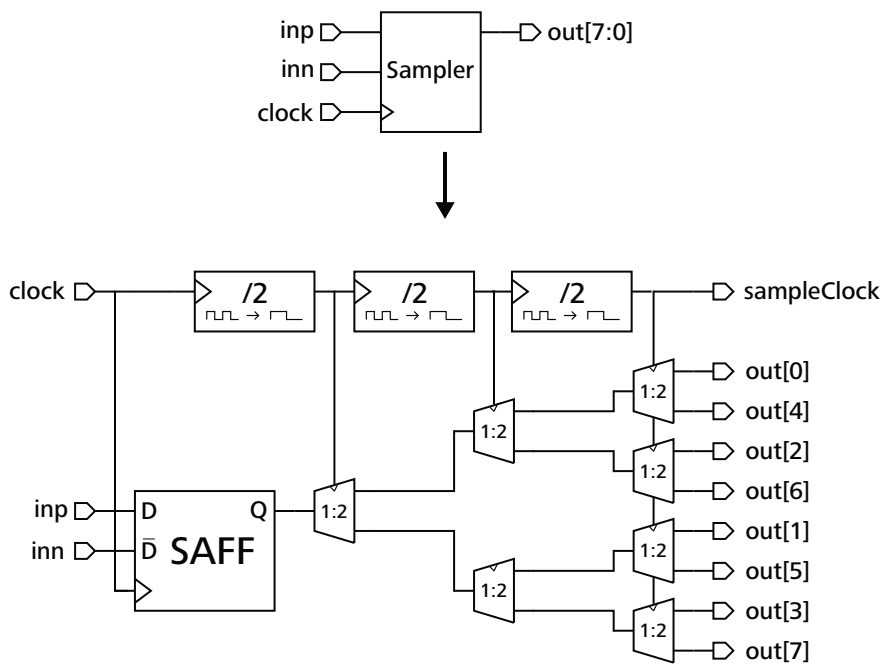Figure 4.12: Sense-Amplifier based Flip-Flop schematic

Figure 4.13: Sampler with 1:8 Deserializer

(a) Deserializer block diagram

(b) Dynamic Flip-Flop schematic (rising edge)

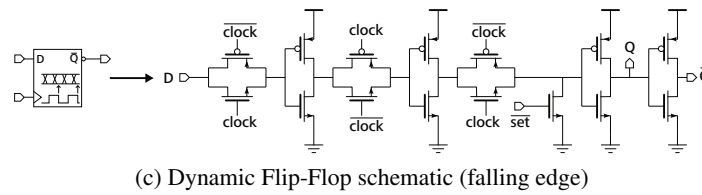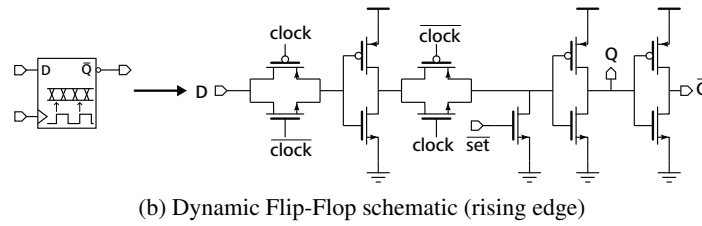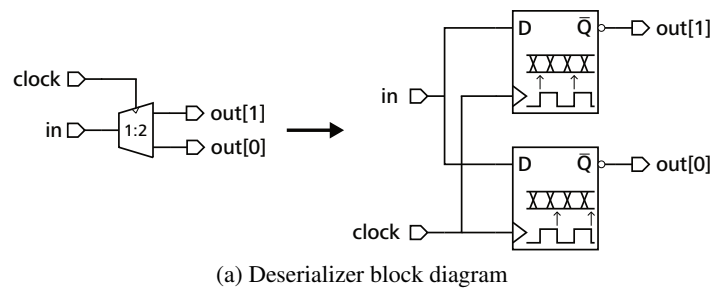(c) Dynamic Flip-Flop schematic (falling edge)

Figure 4.14: 1:2 Deserializer

Samplers, the incoming data is represented by three vectors of 8 bits: *Pre8*, *Edge8* and *Post8*.
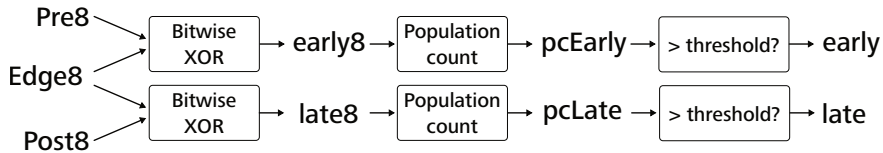

Figure 4.15: Reduction of *early8/late8* signals to *early/late* signal

The incoming 8-bit vectors are processed as depicted in Figure 4.15. This results in the intermediate *early8* and *late8* signals. The population count of the *early8* and *late8* signal determines how many of the 8 sample moments were interpreted as either early or late. The final *early* and *late* signals are therefore directly derived from the *early8* and *late8* signals: if the population count of *early8* and *late8* is higher than the threshold, the *early* and *late* will be high accordingly. To further clarify the population count operation, some examples are given for a threshold of 3 in Table 4.2. This PD Logic implementation results in the single-bit *early* and *late* signals that are required for the Charge Pump.

Input	Population Count	> threshold?
00000100	1	0
11011100	5	1
10010100	3	0
01101111	6	1
10101001	4	1

Table 4.2: Phase Detector Reduction examples (threshold = 3)

The operation of the 10 Gbit/s CDR depends on three samples (Pre, Edge, Post), yet only 2 Samplers are used, as is shown in Figure 4.11. This is possible since the PD Logic stores the previous state, which is then used as the third sample, as shown in Table 4.3.

In the 40 Gbit/s CDR, only part of the incoming data to the PD Logic actually carries information from the incoming data stream, since the samplers

Pre[7:0]	Edge[7:0]	Post[7:0]
$D_{n-1}[15:8]$	$D_n[7:0]$	$D_n[15:8]$

Table 4.3: Phase Detector Data Selection (10/2.5 Gbit/s CDR)

are never enabled all simultaneously. Therefore, the PD Logic picks out the information bits from the incoming data based on the current channel selection as shown in Table 4.4. In the 10/2.5 Gbit/s CDR this selection procedure is not needed, since no sub-sampling operation was implemented there and hence all incoming data carries useful information.

	Pre[7:0]	Edge[7:0]	Post[7:0]
Channel 0	$D_{n-1}[63:56]$	$D_n[7:0]$	$D_n[15:8]$
Channel 1	$D_n[15:8]$	$D_n[23:16]$	$D_n[31:24]$
Channel 2	$D_n[31:24]$	$D_n[39:32]$	$D_n[47:40]$
Channel 3	$D_n[47:40]$	$D_n[55:48]$	$D_n[63:56]$

Table 4.4: Phase Detector Data Selection (40 Gbit/s CDR)

4.2.4   Voltage Controlled Oscillator

The Voltage Controlled Oscillator (VCO) provides the clock that is used to trigger the samplers in the PD. As was explained in Section 4.2.3.1, the implementation of the 40 Gbit/s and the 10/2.5 Gbit/s PD is slightly different. Consequently, the VCOs used in each case must be adapted to the input clock requirements of the different PDs.

This means two distinct VCOs have been implemented: one for the 40 Gbit/s sub-sampling CDR, and one for the 10/2.5 Gbit/s full-rate CDR.

4.2.4.1   Phase Noise

The Phase Noise (PN) of the VCO outside of the CDR bandwidth is a major contributor to the jitter of the recovered clock. As a result, care has been taken to limit the PN of the VCOs in the CABINET CDRs. For the 40 Gbit/s and 10 Gbit/s cases, the CDR bandwidth was chosen at 16 MHz, and therefore the PN is measured at 16 MHz. Based on previous published CDRs [7], a phase noise of -86 dBc/Hz at an offset frequency of 1 MHz suffices for an oscillator running at 5 GHz. Taking into account our oscillator runs at 10 GHz (+6 dB) and the offset frequency is 16 MHz instead of 1 MHz ($4\times$ -6 dB), we arrive at a phase noise specification of maximally -104 dBc/Hz at an offset frequency of 16 MHz. On the other hand, in the 2.5 Gbit/s case, the CDR bandwidth was reduced to 4 MHz and consequently the PN is observed at 4 MHz. Starting from the same -86 dBc/Hz at 1 MHz, the oscillator now runs at 2.5 GHz (-6 dB) and the offset frequency is increased to 4 MHz ($2\times$ -6 dB), resulting in the same phase noise upper limit of -104 dBc/Hz.

4.2.4.2   VCO Architecture

Today, two main oscillator architectures are in use: LC oscillators and ring oscillators. LC oscillators use a resonant LC-tank, while a ring oscillator consists of a loop of delay cells that satisfies the Barkhausen oscillation criteria [10–12].

LC oscillators typically achieve lower phase noise, but are relatively large due to the inductor of the LC tank. Furthermore, they are known for their limited tuning range. On the other hand, ring oscillators consume only a small area and typically have large tuning ranges. Unfortunately, this is generally accompanied by a higher phase noise.

Since the phase noise requirements for CDRs are not too demanding, ring oscillators are a good choice for this application. Furthermore, ring oscillators can easily provide multiple clock phases, which is beneficial for the 40 Gbit/s sub-sampling CDR.

As a result, the VCOs for the CDRs in the CABINET were implemented as ring oscillators.

4.2.4.3   40 GHz Voltage-Controlled Oscillator

The 40 Gbit/s PD makes use of a series of eight 10 Gbit/s samplers instead of three 40 Gbit/s samplers, which would be overly power consuming, especially when considering 75% of the recovered bits would have to be discarded. However, to trigger this series of 8 samplers, 8 clock phases at 10 GHz are required. This means the output of the 40 GHz VCO is a slower 10 GHz, but needs to provide 8 equally spaced clock phases at this frequency.

The requirement of 8 clock phases for the 40 Gbit/s PD operation immediately fixes the 40 Gbit/s VCO architecture to a 4-stage ring oscillator with 4 differential delay cells, as shown in Figure 4.16.
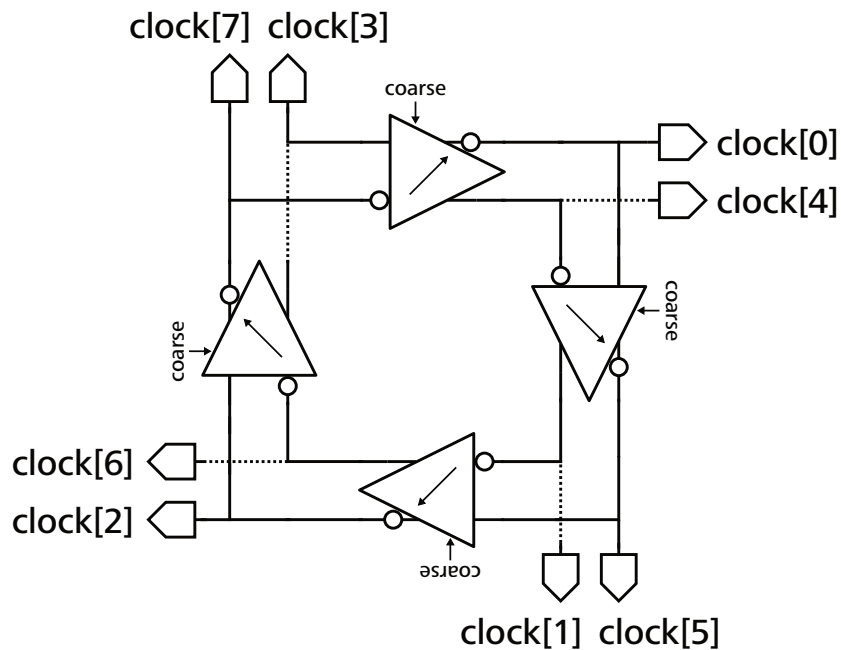

Figure 4.16: 40 GHz Voltage-Controlled Oscillator

The delay cell used in the 40 Gbit/s VCO is shown in Figure 4.17. Due to the limited supply voltage tolerated in the used 40 nm CMOS technology, tail currents were avoided to maximize the output voltage swing. A cross-coupled common source amplifier with resistive load was used to avoid the input capacitance of the PMOS, which is present in the typically used CMOS inverters. This helps to maximize the oscillation frequency.
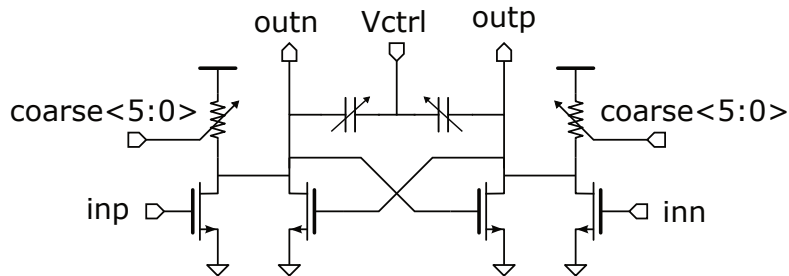

Figure 4.17: 40 GHz VCO Delay Cell

Coarse tuning of the delay cell is accomplished by a tunable resistive load, whose value is set by digital control bits, while fine tuning is achieved by means of a varactor tuned by a control voltage. Both mechanisms are indicated on Figure 4.17.

The physical implementation of the 40 GHz VCO is shown in Figure 4.18, with the four differential delay cells highlighted. The dimensions of this VCO are approximately $106\,\mu$m $\times$ $91\,\mu$m, resulting in an area of only $0.009649\,\text{mm}^2$.

Since it is very important to have the output phases of this VCO evenly spaced, special care has been taken to avoid any influence of the layout on the phase-spacing. Consequently, the interconnection of the delay cells has been carefully matched in length, to avoid any delay variations these interconnections have between the output phases due to wiring.

Furthermore, on the left and right side of the delay cells, one can see half-cell dummy instances (Figure 4.18). The dummy cells are used to keep the physical environment identical for all delays cells. Doing so helps to reduce manufacturing variations across different delay cells, which would influence the spacing between the output phases.

The 40 GHz VCO was designed to output 8 phases of 10 GHz, exhibit a VCO gain of about -350 MHz/V and phase noise of maximum -104 dBc/Hz at 16 MHz, all with a power consumption of approximately 30 mW. These
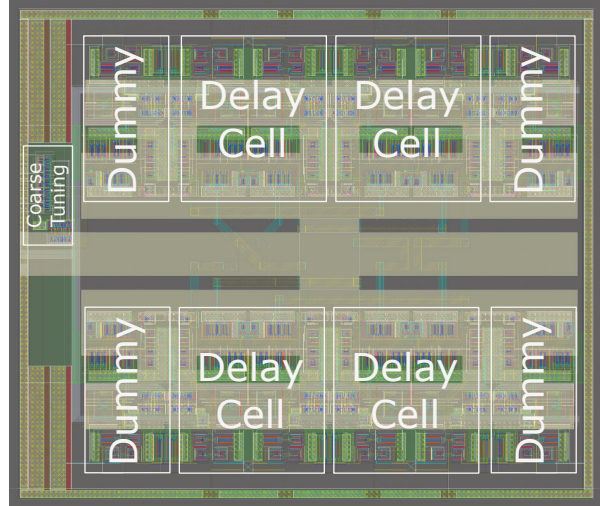
Figure 4.18: 40 GHz VCO Delay Cell Physical Implementation

specifications were verified by post-layout simulations across process and temperature variation corners, as shown in Table 4.5.

Specification	40 GHz VCO Value
Oscillation Frequency	8×10 GHz
VCO Gain K_{VCO}	-343.08 MHz/V
Maximum Phase Noise	-112.8 dBc/Hz @ 16 MHz
Power consumption	30.42 mW

Table 4.5: Post-layout simulation performance of 40 GHz VCO

Additionally, Figure 4.19a presents the simulated coarse tuning range for V_{ctrl} = 0.55 V, for the Slow-Slow process corner at 80°C (SS80) and in the Fast-Fast process corner at 0°C (FF0). It is clear the desired 10 GHz oscillation frequency (of which there are 8 phases) is achievable with the designed VCO.

The VCO control voltage tuning characteristic is depicted in Figure 4.19b for the coarse setting corresponding to the 10 GHz oscillation frequency. The control voltage V_{ctrl} is swept from 0.35 V to 0.75 V and the resulting oscillation frequency is displayed. This tuning characteristic also exposes
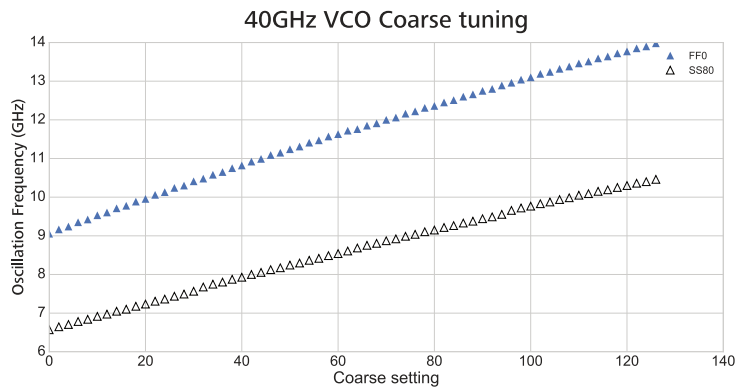
(a) 40 GHz VCO Coarse Tuning Simulation
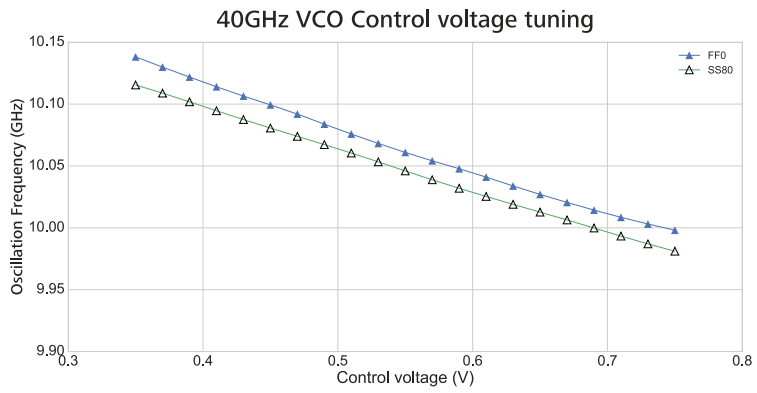
(b) 40 GHz VCO Control Voltage Tuning Simulation

Figure 4.19: 40 GHz VCO Simulation

the VCO gain K_{VCO}, which is on average -343.08 MHz/V and has limited variation over process and temperature.

4.2.4.4   10/2.5 GHz Voltage-Controlled Oscillator

Since the 10/2.5 Gbit/s CDR uses full-rate samplers and needs only 2 clock phases, the VCO architecture is simplified to a 3-stage, single-ended architecture. To provide a differential output in the case of 10 Gbit/s operation, two such 3-stage rings are coupled, as is shown in Figure 4.20a. To reduce power consumption for 2.5 Gbit/s operation, only one of the two 3-stage rings is enabled and the coupling of the two rings is disabled. The core VCO still oscillates at 10 GHz, which is divided by four to output 2.5 GHz. This divider also provides the differential output. The VCO configuration for the 2.5 Gbit/s operation is shown in Figure 4.20b.

(a) 10 GHz Voltage-Controlled Oscillator

(b) 2.5 GHz Voltage-Controlled Oscillator
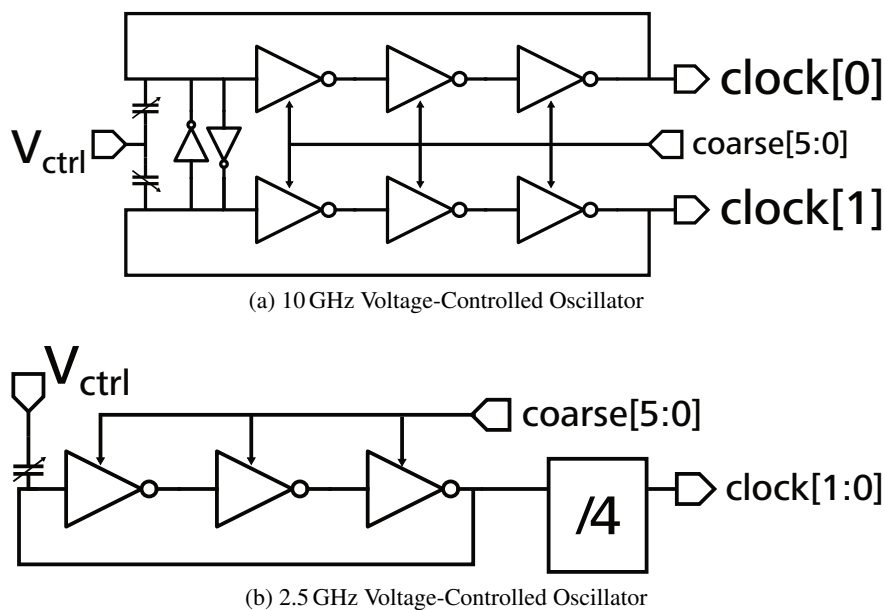
Figure 4.20: 10/2.5 GHz Voltage-Controlled Oscillator

The delay cell used for the 10/2.5 GHz VCO is shown in Figure 4.21. Like in the 40 GHz VCO case, it is a resistively loaded common source amplifier. Furthermore, the tuning mechanism is very comparable, providing coarse tuning by means of a tunable resistive load.

Contrary to the 40 GHz VCO, the fine tuning of the delay cell is not imple-
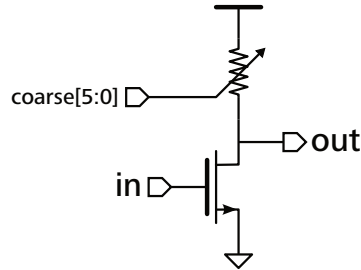
Figure 4.21: 10 GHz VCO Delay Cell

mented on the delay cell level, but on the level of the core VCO, as indicated on Figure 4.20a and Figure 4.20b. This is possible because only one of the three generated phases is used, which means the delay of the used cells can vary with respect to each other, as long as the total delay of the ring corresponds to the desired oscillation frequency. Implementing the fine tuning on the core VCO level offers the advantage that the used varactor can be bigger, which leads to better manufacturability and less variations on the varactor. Moreover, this simplifies the layout of the VCO, since the control voltage must only be routed to one varactor instead of three.

The layout of the 10/2.5 GHz VCO is shown in Figure 4.22. To support both rates, the same VCO is used, but in a different configuration. The physical dimensions of this VCO amount to a height of 58 μm and a width of 91 μm, resulting in an area of only 0.005369 mm^2, which is only half of the 40 GHz VCO.

Figure 4.22a represents the layout configuration for the 10 GHz VCO, with the 3 differential delay cells highlighted. On the left of the delay cells, the tuning varactor is shown. Since only one output phase is used in this configuration, the spacing between phases is not as critical. Therefore, no dummy cells have been used in this layout to save area. However, since a differential output is desired, the complete VCO is made symmetrical across the X-axis.

When operating at 2.5 GHz, the active configuration is as shown in Figure 4.22b. Since one of the two loops has been shut down to reduce power consumption, the single-ended delay cells have been highlighted here. Additionally, the divide-by-4 used to reduce the output frequency of the VCO to the desired range is indicated.

In the 10 GHz configuration, the VCO was designed to output a differential 10 GHz signal with a VCO gain of about -350 MHz/V and a phase noise
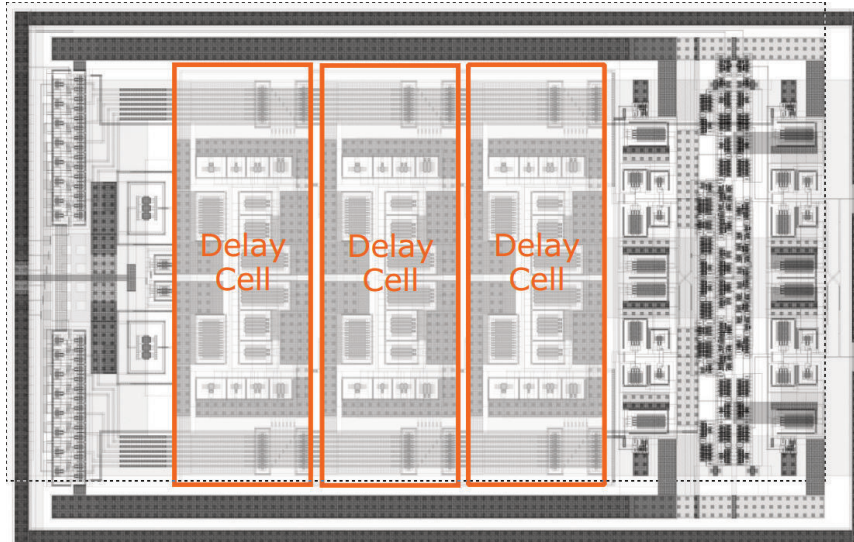
(a) Configuration for 10 GHz VCO
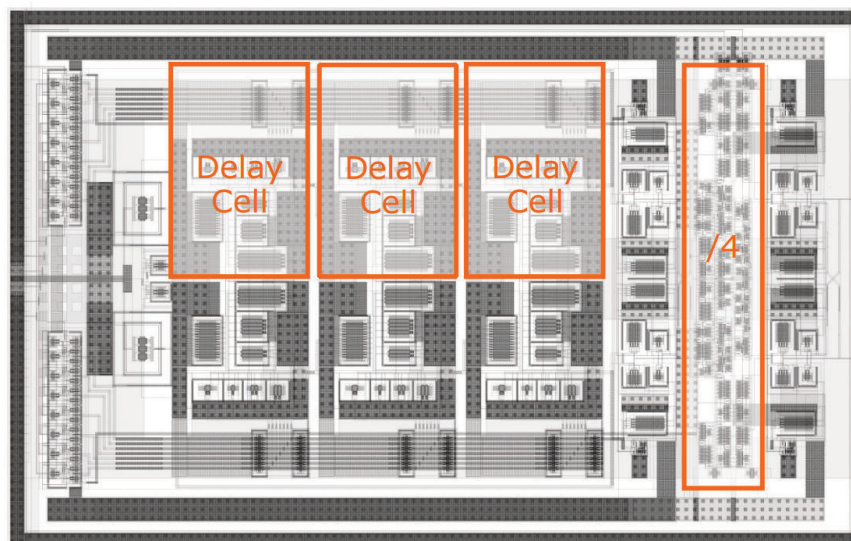
(b) Configuration for 2.5 GHz VCO

Figure 4.22: Layout of 10/2.5 GHz VCO: same block in different
configurations

of no more than -104 dBc/Hz at 16 MHz. In the 2.5 GHz configuration, the same core VCO ring is used, but fed through a divide-by-4 block. This means the VCO gain is only -87.5 MHz/V. The phase noise at 4 MHz from the carrier should not exceed -104 dBc/Hz. Power consumption of the VCO is designed as maximally 10 mW in the 10 GHz configuration, and is reduced to less than 5 mW in the 2.5 GHz configuration.

Like for the 40 GHz VCO case, these numbers were verified by post-layout simulations across process and temperature variation corners as shown in Table 4.6.

Specification	10 GHz VCO Value	2.5 GHz VCO Value
Oscillation Frequency	$2\times$ 10 GHz	$2\times$ 2.5 GHz
VCO Gain K_{VCO}	-345.16 MHz/V	-84.74 MHz/V
Phase Noise	-109.5 dBc/Hz	-106 dBc/Hz
PN Frequency Offset	16 MHz	4 MHz
Power consumption	9.672 mW	4.978 mW

Table 4.6: Post-layout simulation performance of 10/2.5 GHz VCO

The tuning ranges of the 10 GHz and 2.5 GHz VCO were simulated and are shown in Figure 4.23 and Figure 4.24 respectively. The coarse tuning was simulated with V_{ctrl} = 0.55 V in the Slow-Slow process corner at 80°C and in the Fast-Fast process corner at 0°C (Figure 4.23a and Figure 4.24a). In both cases the desired oscillation frequency (10 GHz and 2.5 GHz) is achievable. Furthermore, at the coarse setting that corresponds to the desired oscillation frequency, the control voltage V_{ctrl} is simulated from 0.35 V to 0.75 V, revealing the VCO gain K_{VCO} and the tuning range within one coarse setting (Figure 4.23b and Figure 4.24b).

4.2.5   Charge Pump and Loop Filter Sizing

The performance of the Bang-Bang CDR is dependent on the combination of the Charge Pump current, Loop Filter sizing and the VCO gain. The VCO gain is typically not very flexible, as it is largely dictated by the VCO architecture and the used technology. Therefore, a bang-bang CDR is typically designed with 3 degrees of freedom: the Charge Pump current, the Loop Filter resistor and the Loop Filter capacitor.
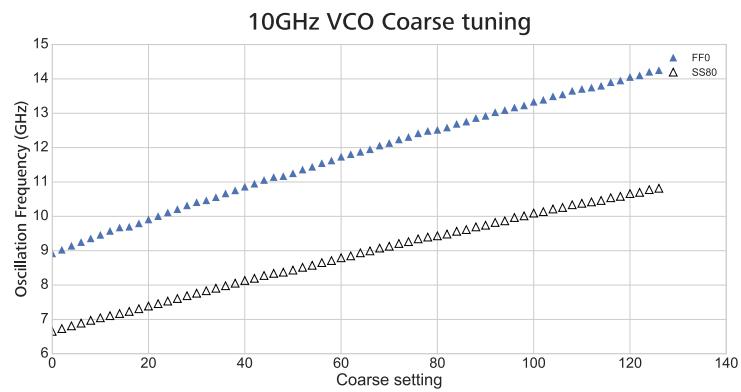
(a) 10 GHz VCO Coarse Tuning Simulation
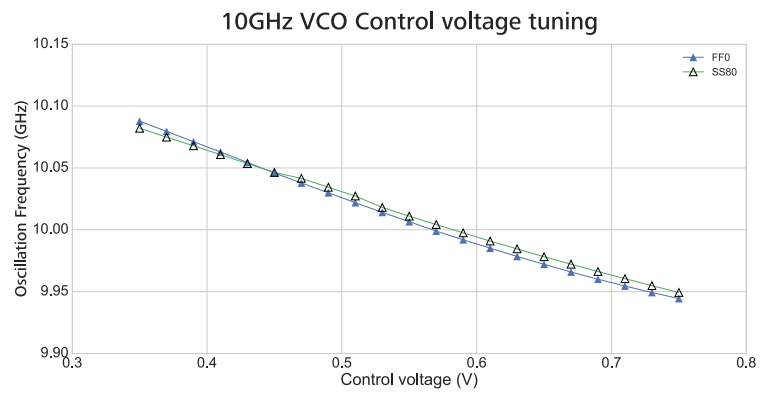
(b) 10 GHz VCO Control Voltage Tuning Simulation

Figure 4.23: 10 GHz VCO Simulation

(a) 2.5 GHz VCO Coarse Tuning Simulation

(b) 2.5 GHz VCO Control Voltage Tuning Simulation

Figure 4.24: 2.5 GHz VCO Simulation

Thanks to the availability of a linear model, the analysis of a Linear CDR is straightforward. The analysis of a Bang-Bang CDR is however quite complicate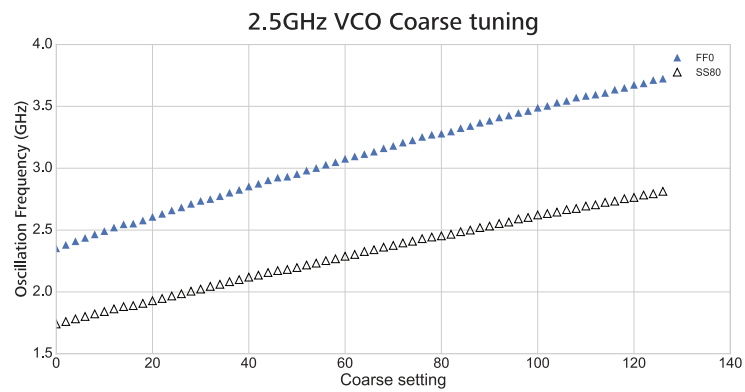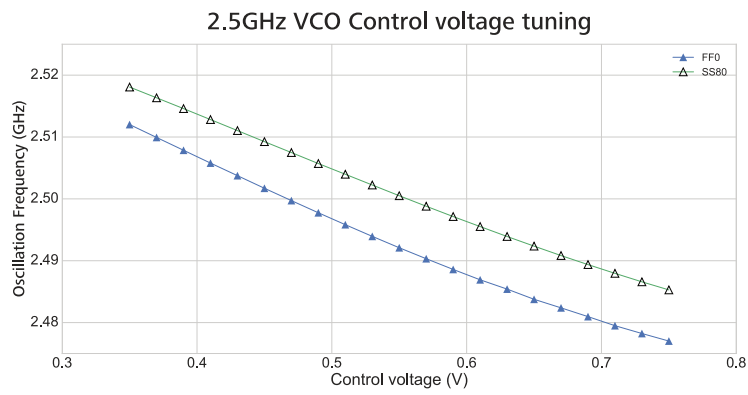d due to the non-linear Phase Detector. As a result, a great deal of research has been conducted in an effort to provide an accurate yet simplified analysis [13–24]. One approach is to linearize the BB-PD gain, which allows to analyze the loop as if it were a Linear CDR. It is clear this path greatly simplifies the design of the CDR and was therefore followed for the design of the CABINET CDR.

The sizing of the Charge Pump and the Loop Filter is based on the Linear CDR analysis performed in [25], where the PD gain, K_{PD}, of the Linear PD is replaced by the linearized gain of the BB-PD. This linearization is the outcome of the analysis with Describing Functions that was presented in [22] and results in a K_{PD} as shown in Equation 4.2, where A is the amplitude of the input phase.

$$K_{PD} = \frac{4}{\pi} \frac{1}{|A|} \tag{4.2}$$

The Loop Filter used in the CABINET CDR is a second-order filter consisting of one resistor R_{filter} and two capacitors C_{filter} and C_{aux} connected as shown in Figure 4.25. The auxiliary capacitor C_{aux} is typically added to improve the transient behavior of the system by filtering out undesired high frequency signals on the control voltage V_{ctrl}. The transfer function of the Loop Filter is given in Equation 4.3, under the assumption that the auxiliary capacitor C_{aux} is much smaller than the filter capacitor C_{filter}. This results in the open loop transfer function $H_{open}(s)$ of the CDR given by Equation 4.8.
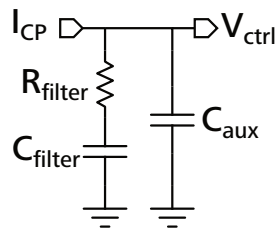

Figure 4.25: Loop Filter Architecture

$$LF(s) = \frac{1}{sC_{filter}} \cdot \frac{s/\omega_z + 1}{s/\omega_p + 1} \tag{4.3}$$

$$\omega_z = \frac{1}{RC_{filter}} \tag{4.4}$$

$$\omega_p = \frac{1}{RC_{aux}} \tag{4.5}$$

$$H_{open} = K_{PD} \cdot I_{CP} \cdot LF(s) \cdot \frac{K_{VCO}}{s} \tag{4.6}$$

$$\Leftrightarrow H_{open} = \frac{4}{\pi} \frac{1}{|A|} \cdot I_{CP} \cdot \frac{1}{sC_{filter}} \cdot \frac{s/\omega_z + 1}{s/\omega_p + 1} \cdot \frac{K_{VCO}}{s} \tag{4.7}$$

$$\Leftrightarrow H_{open}(s) = \frac{4}{\pi} \frac{1}{|A|} \frac{I_{CP}K_{VCO}}{C_{filter}} \cdot \frac{1}{s^2} \cdot \frac{s/\omega_z + 1}{s/\omega_p + 1} \tag{4.8}$$

As was described in [25], the pole-zero ratio $\frac{\omega_p}{\omega_z}$ determines the trade-off between speed and stability: a higher ratio translate to a higher phase margin in the open loop transfer function, and therefore a more stable system. To prevent excessive overshoot, a phase margin of at least 60°is desired. However, to further reduce the jitter transfer peaking the phase margin is chosen to be 85°, which corresponds to a pole-zero ratio of about 1000 [25]. This requirement fixes the relation between C_{filter} and C_{aux} (Equation 4.9).

$$C_{filter} = 1000 \times C_{aux} \tag{4.9}$$

Furthermore, the cut-off frequency ω_c of the closed-loop transfer function is found from the pole-zero pair as shown in Equation 4.10. From this relation, both ω_z and ω_p can be defined in terms of ω_c. From the definition of ω_z, the value of R_{filter} is then found using Equation 4.13. Note that the cut-off frequency ω_c is not fixed, but dependent on the CDR selection: 16 MHz for the 40 and 10 Gbit/s CDR, and 4 MHz for the 2.5 Gbit/s CDR.

$$\omega_c = \sqrt{\omega_p \cdot \omega_z} \tag{4.10}$$

$$\Rightarrow \omega_z = \frac{\omega_c}{10\sqrt{10}} \tag{4.11}$$

$$\Rightarrow \omega_p = 10\sqrt{10} \times \omega_c \tag{4.12}$$

$$R_{filter} = \frac{1}{\omega_z C_{filter}} \tag{4.13}$$

The value for the filter capacitor C_{filter} is found from the requirement that the magnitude of the open-loop transfer function is 1 at ω_c.

$$|H_{open}(\omega_c)| = |\frac{4I_{CP}K_{VCO}}{\pi|A|C_{filter}} \cdot \frac{1}{-\omega_c^2} \cdot \frac{\omega_c/\omega_z + 1}{\omega_c/\omega_p + 1}| = 1 \qquad (4.14)$$

$$\Leftrightarrow C_{filter} = 10\sqrt{10} \cdot \frac{4}{\pi}\frac{1}{|A|}\frac{I_{CP}K_{VCO}}{\omega_c^2} \qquad (4.15)$$

A represents the amplitude of the input jitter. The value of A used for the design of the CDR is empirical and chosen as 0.15 Unit Interval (UI), a value also used in literature [21], which is based on the input jitter mask in typical ITU-T jitter tolerance specifications. The only unknown left is the Charge Pump current I_{CP}, which is dimensioned sufficiently low to limit the power consumption, but high enough to yield realistic capacitor values.

The CABINET supports multiple rates: 40 Gbit/s, 10 Gbit/s and 2.5 Gbit/s. The CABINET CDR requirements for 40 Gbit/s and 10 Gbit/s operation are quite similar, due to the inherent 1:4 sub-sampling nature during 40 Gbit/s operation. In that configuration, the recovered clock is effectively 10 GHz, as is the case during 10 Gbit/s operation. Since the VCOs in the 40 Gbit/s CDR and the 10 Gbit/s CDR oscillate at the same frequency and have the same VCO gain, both CDRs can be treated similarly. Therefore, while both cases are not exactly the same, it was decided to use the same CDR configuration for both to simplify the design. The CDRs remain physically separate entities, but the Charge Pump and Loop Filter are implemented equally in each CDR.

When comparing the 2.5 Gbit/s CDR to the 10 Gbit/s one, the VCO output is only 2.5 GHz and the VCO gain is divided by four. Therefore, to support 2.5 Gbit/s operation, the CABINET CDR configuration must change. Furthermore, the closed-loop CDR bandwidth was chosen 16 MHz for 40 Gbit/s and 10 Gbit/s operation, but was reduced to 4 MHz for 2.5 Gbit/s operation.

The auxiliary capacitor C_{aux} is implemented on-chip: this allows it to compensate the inductive bondwires that are used to connect the off-chip C_{filter}. Furthermore, its presence helps to reduce high-frequency noise. The low value of C_{aux} also means it is difficult to make it tunable. Therefore, it was chosen to make C_{aux} a fixed value. The C_{filter} is implemented off-chip due to its large value, which would take up much of the expensive die area. Its value is more or less fixed, since C_{aux} has a fixed value. Configuration of the filter is therefore accomplished by means of the resistor R_{filter} which is implemented as a tunable component.

When changing rates, both K_{VCO} and ω_c change, which can not be compensated for only by the tunable resistor. As a result, the Charge Pump current I_{CP} is also implemented as a tunable parameter, allowing us to tune R_{filter} and I_{CP} to their desired values according to the design equations we derived previously. The nominal values for the different CDR components are summarized in Table 4.7. Note that the implemented tunability also allows to adjust ω_c in function of the input jitter $|A|$.

	40/10 Gbit/s	2.5 Gbit/s
ω_c	16 MHz	4 MHz
K_{VCO}	≈-345 MHz/V	≈-85 MHz/V
I_{CP}	200 μA	60 μA
R_{filter}	540 Ω	1.83 kΩ
C_{filter}	582 pF	685 pF
C_{aux}	≈700 fF	≈700 fF

Table 4.7: CDR Design component values

4.2.5.1 Charge Pump implementation

The Charge Pump implementation is depicted in Figure 4.26. The *up* and *down* signals are coming from either the Frequency Detector or the Phase Detector, depending on the *locked* signal of the FLL. These drive transistors that act as switches and either switch the current to the *out* node, or to a dummy node. The implementation with a dummy node was chosen to increase the switching speed of the Charge Pump, since the *up* and *down* pulses could have repetition rates up to 1.25 GHz.

The reference currents, both on the top as in the bottom, are equal and are scaled copies from the input current. The scaling factor is tunable, making the Charge Pump current configurable from 10 μA to 200 μA. As was mentioned in Section 4.2.5, the current has to be tunable to support multiple rates. The range of the implemented Charge Pump exceeds the range required by the nominal CDR design. However, it was decided a somewhat larger tuning range would be favorable to compensate for possible manufacturing variations in the ASIC.
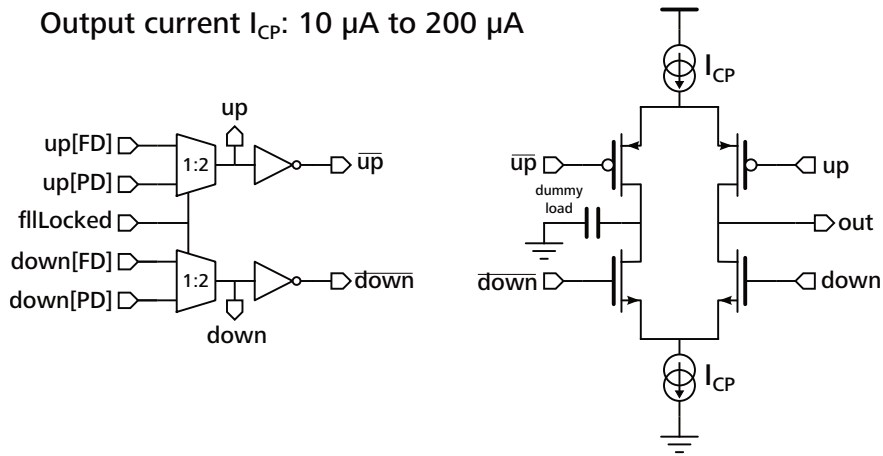
Output current I$_{CP}$: 10 µA to 200 µA

Figure 4.26: Charge Pump Implementation

4.2.5.2   Loop Filter implementation

As described in Section 4.2.5 the on-chip Loop Filter mainly consists of a resistor R_{filter} that is tunable in the range of 0 to 12.6 kΩ, once more allowing for some variation compensation. The filter capacitor C_{filter} is placed off-chip to allow for large capacitance values, since these typically consume a lot of expensive chip area. Furthermore, this leaves the possibility to change the capacitor value after fabrication. The auxiliary capacitor C_{aux} is implemented on-chip and has a value of about 700 fF. The implementation is visualized in Figure 4.27.
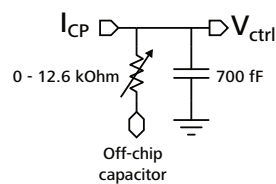
Figure 4.27: Loop Filter Implementation

4.2.6   CDR Locking Behavior

To verify that the design of the CDR results in a phase-lock, simulations were performed to see if the control voltage *Vctrl* settles to a steady state

value, while locking correctly to the incoming datastream. The result of these simulations is shown in Figure 4.28 for the 40 Gbit/s and 10 Gbit/s case, and in Figure 4.29 for the 2.5 Gbit/s case. In both cases, the simulation confirms the correct operation of the CDR loop: the control voltage reaches a steady state value, meaning the CDR is locked, and once this lock is achieved the *error* signal stays low (Figure 4.28b and Figure 4.29b), meaning the received bits are correct.
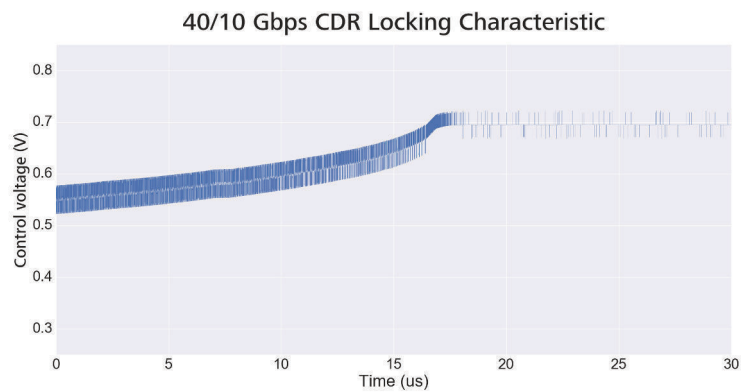
(a) Control voltage locking characteristic
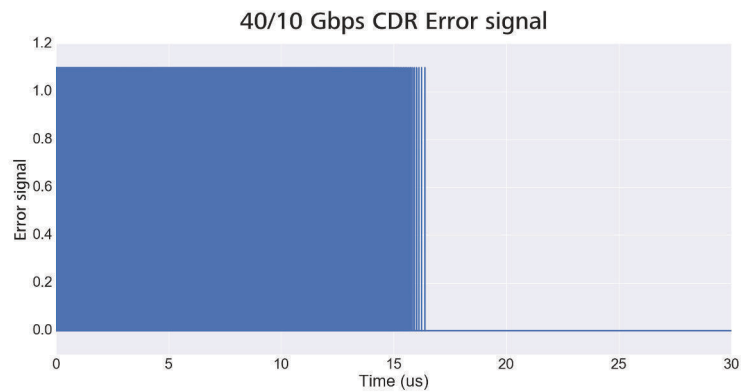
(b) Error signal (0: received bit correct, 1.1: received bit incorrect)

Figure 4.28: 40 Gbit/s and 10 Gbit/s CDR Simulated Operation

(a) Control voltage locking characteristic

(b) Error signal (0: received bit correct, 1.1: received bit incorrect)

Figure 4.29: 2.5 Gbit/s CDR Simulated Operation

4.3   MAC Preprocessor

The MAC Preprocessor is a digital block that essentially performs two major tasks: (1) extract the CABINET configuration from the CBi Header, and (2) extract and process the payload of the CBi Frame to either simply repeat to the lower CBi Level or deliver to the FPGA.
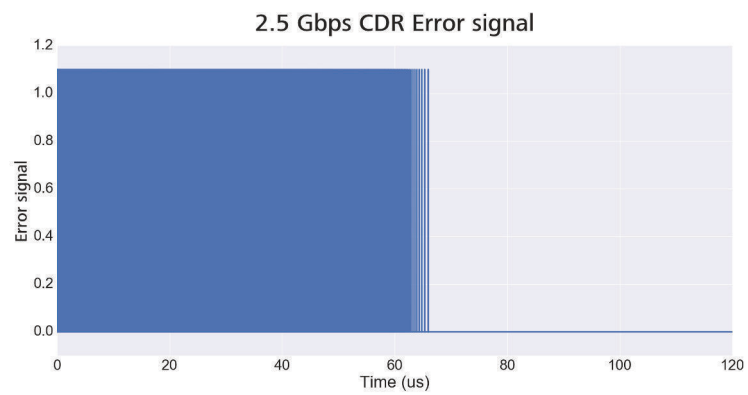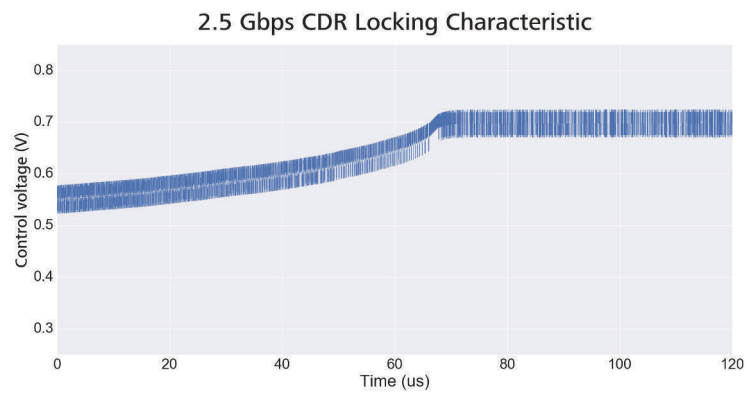
Before the MAC Preprocessor can initiate processing, the CDR must be locked and correctly recovering the incoming data and the related clock. When CDR lock is achieved, the first step is a synchronization procedure, which is common to both CBi Repeaters and CBi End-ONTs. The following step is dependent on the configuration of the receiver (i.e. CBi Repeater or CBi End-ONT mode).

4.3.1   Synchronization procedure

The synchronization procedure starts with a search for the SYNC word in the CBi Header, during which the procedure is said to be in the *hunt* state. Once the SYNC word has been found, the CBi Header is parsed and the Repeater/eNd-ont IDentifier (RNID) is read. Based on the received RNID and the configured RNID of the receiving CABINET, the CBi Header lane offset is calculated and applied in the *pre-sync* state, resulting in correct sampling of the CBi Header for this particular receiver.

Once the SYNC pattern is verified and the RNIDs match, the *sync* state is reached. The synchronization procedure then keeps monitoring the incoming CBi Frames to continuously verify the SYNC pattern and RNID. To allow for transmission errors, up to 3 verification fails are tolerated, moving to the *re-sync* state. However, a fourth verification fail results in a restart with a clean slate. The complete synchronization procedure is shown in Figure 4.30.

4.3.2   Repeater Mode

In Repeater mode, the CBi Header is descrambled to obtain the bandwidth map (BWMAP), which is then parsed to get the necessary parameters regarding the line rate. The payload of incoming CBi Frames is then forwarded to the output without any additional processing. For every next CBi Frame received, the BWMAP is monitored for changes. If the DSBWMAP flag is set, the configuration of the receiver is adjusted accordingly and the receiver waits for a new CBi Frame.

Synchronization

Hunt

SYNC pattern match
(adjust lane)

SYNC pattern and/or
RNID verification fail

Pre-Sync

SYNC pattern and
RNID verified

Sync

Both verified

SYNC pattern and/or
RNID verification fail

SYNC pattern and
RNID verified (once)

Re-Sync

SYNC pattern and/or
RNID verification fail

M-1 consecutive SYNC pattern
and/or RNID verification fails (M=3)
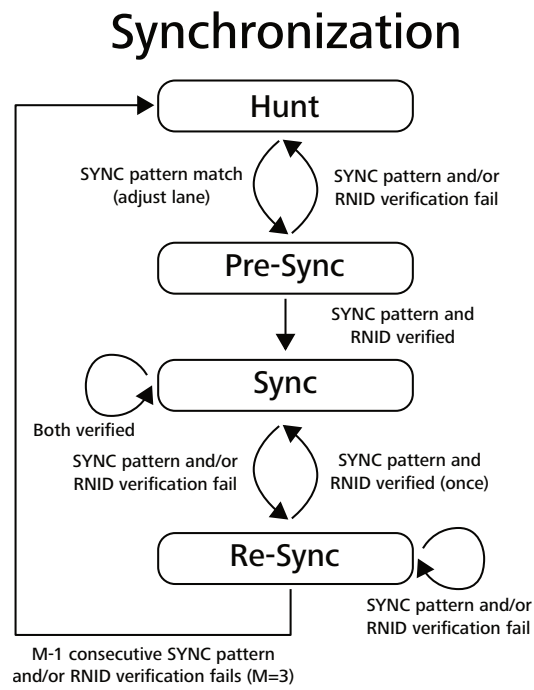
Figure 4.30: MAC Preprocessor Synchronization Procedure

Repeater

```
Wait for CBi Frame
        ↓
  Synchronization
        ↓
Descramble CBi Header
   Obtain BWMAP
        ↓
   Parse BWMAP
Obtain DSBWMAP (flag, rate)
        ↓
    DSBWMAP
    Flag Set?  ── Yes
        │
        No
        ↓
  Forward payload
```

End-ONT

```
Wait for CBi Frame
        ↓
  Synchronization
        ↓
Descramble CBi Header
   Obtain BWMAP
        ↓
   Parse BWMAP
Obtain DSBWMAP (flag, rate, offset)
        ↓
Select payload clock
 According to BWMAP rate
        ↓
 Adjust CDR Phase
Determined by BWMAP offset
        ↓
Retrieve payload data
Sample output with payload clock
        ↓
 Descramble payload
        ↓
Output data and clock
```
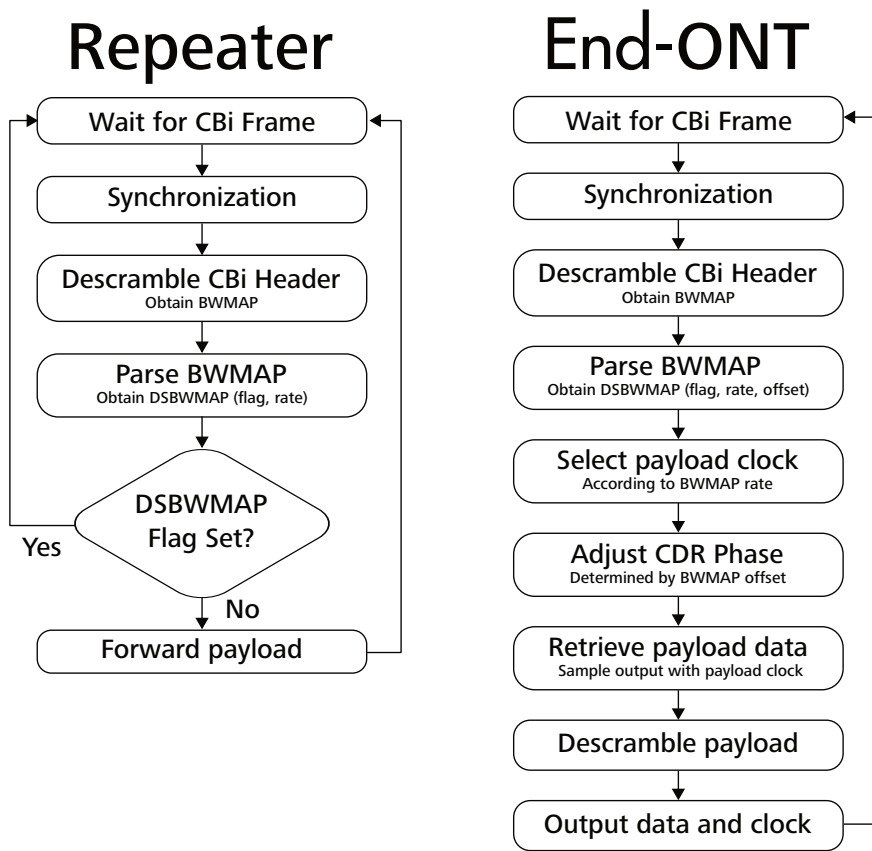
Figure 4.31: MAC Preprocessor Frame Processing

Reducing the processing in the CBi Repeaters as much as possible by only parsing the CBi Headers and simply forwarding the payload after a decimation operation reduces the power consumption of a CBi Repeater.

4.3.3   End-ONT Mode

In End-ONT mode, like in Repeater mode, the BWMAP is retrieved from the descrambled CBi Header data. The line rate from the BWMAP determines the selection of the correct payload clock.

When operating at 40 Gbit/s, the CDR is inherently sub-sampling. Therefore, the CDR sampling channel is adjusted to the correct channel if necessary: this is determined by the offset value in the BWMAP.

Subsequently, the payload clock is used to subsample the payload data, which is then descrambled and sent to the FPGA.

4.4   Analog Back-End

Since the MAC Preprocessor is implemented by use of the provided Standard Cell Library, its maximum frequency of operation is limited to 1.25 GHz. However, a CBi Repeater connecting a CBi Level L1 (40 Gbit/s) to a CBi Level L2 (10 Gbit/s) requires an output of 10 Gbit/s. Therefore, when the payload is repeated by the CABINET, the 8×1.25 Gbit/s should be serialized before being sent to the single-lane optical transmitter module, which is the objective of the Analog Back-End.
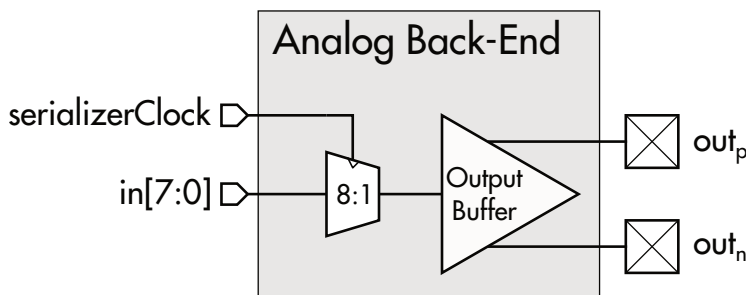

Figure 4.32: Analog Back-End

The Analog Back-End therefore consists of an 8:1 serializer and a high speed output buffer, as shown in Figure 4.32.
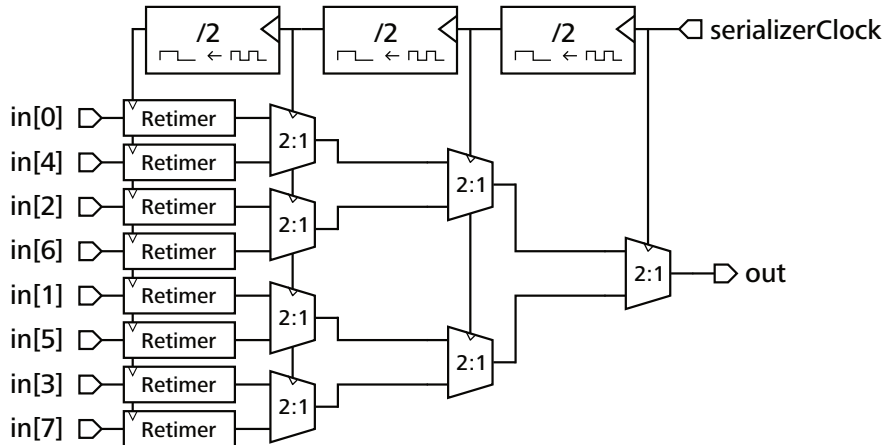
4.4.1   8:1 Serializer

Figure 4.33: 8:1 Serializer

The 8:1 serializer is implemented as a tree of 2:1 serializers (Figure 4.33). The highest frequency clock is the *serializerClock*, which is divided down to provide the lower frequency clocks for the intermediate serializer stages. Since the data comes from the MAC preprocessor, whose clock is not necessarily synchronous with the serializer clock, all datastreams are retimed before entering the 2:1 serializers.

Ideally, a CDR would tackle this asynchronous connection, but for simplicity and optimization of design time, the retimers are configurable to retime on either the rising or the falling edge of the clock. This means if there is a setup or hold time violation by the incoming data, this can be solved manually by retiming on the opposite edge, while a CDR would do this automatically.

4.4.2   Output Buffer

The output buffer is a standard Current-Mode Logic (CML) output driver, consisting of a differential pair that is matched to $50\,\Omega$. A bandwidth of around $7\,\text{GHz}$ suffices, since the highest output rate is $10\,\text{Gbit/s}$. However, the output should be $2.5\,\text{V}$ LVPECL compatible, requiring slower $2.5\,\text{V}$ tolerant devices, which complicates the design of the output buffer. As a result, to achieve the required specifications, the output buffer has a simulated power consumption of about $43\,\text{mW}$, making it a significant contributor to the total power consumption of the CABINET.
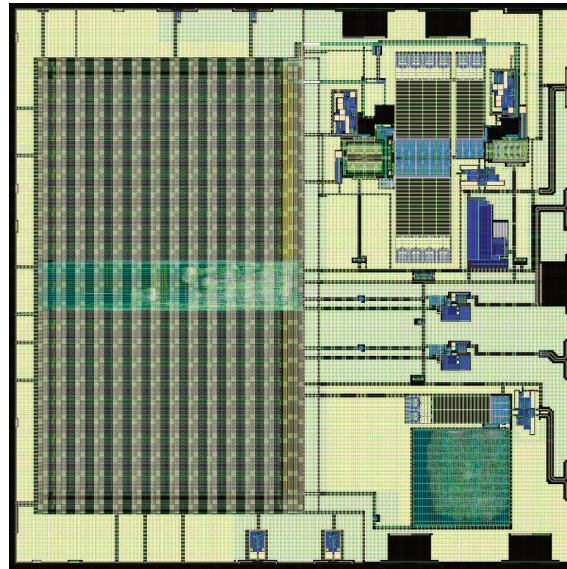
4.5   CABINET ASIC Layout

The CABINET ASIC was manufactured in a 40 nm CMOS technology. The drawn layout is shown in Figure 4.34a, while a photograph of the manufactured CABINET ASIC is shown in Figure 4.34b. The complete chip dimensions measure about 1.85 mm × 1.85 mm.

4.6   Conclusion

This chapter discussed the design and implementation of the CABINET ASIC, which serves as a generic CBi Device, configurable in both Repeater and End-ONT Mode, and supporting multiple rates: 40 Gbit/s, 10 Gbit/s and 2.5 Gbit/s. In the first part, the system architecture of the CABINET ASIC was introduced, revealing that the CABINET consists of three main parts: the Analog Front-End, the MAC preprocessor and the Analog Back-End.

The second part presented the details of the design and implementation of the Analog Front-End, discussing the CDR topology selection, the implementation of the critical building blocks and the Phase Detector Logic algorithms. Additionally, an elaborate explanation of the Voltage-Controlled Oscillators was given. The sizing of the Loop Filter was discussed and simulations were shown confirming that the CDRs lock.

Subsequently, the MAC preprocessor operation was clarified in the third part, presenting the Synchronization Procedure and showing the details of the MAC preprocessor Frame Processing procedure. The fourth part concluded the chapter elaborating on the implementation details of the Analog Back-End. Finally, the CABINET ASIC layout was shown.

(a) Drawn layout of complete CABINET ASIC
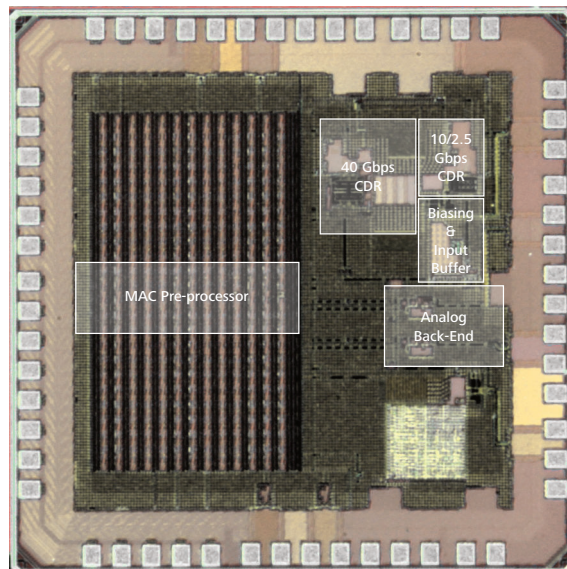
(b) Photograph of manufactured CABINET ASIC

Figure 4.34: Complete CABINET ASIC: drawn layout versus
manufactured ASIC

References

[1] Maxim Integrated. NRZ Bandwidth - HF Cutoff vs. SNR (HFAN-09.0.1 Rev2;04/08). Application note, Maxim Integrated, 2008.

[2] Ming-ta Hsieh and G Sobelman. Architectures for multi-gigabit wire-linked clock and data recovery. *Circuits and Systems Magazine, IEEE*, 8(4):45–57, 2008.

[3] C. R. Hogge. A self correcting clock recovery circuit. *IEEE Transactions on Electron Devices*, 32(12):2704–2706, Dec 1985.

[4] J. D. H. Alexander. Clock recovery from random binary signals. *Electronics Letters*, 11(22):541–542, October 1975.

[5] Behzad Razavi. Challenges in the design of high-speed clock and data recovery circuits. *IEEE Communications magazine*, 40(8):94–101, 2002.

[6] Adrian Fox. Ask the Applications Engineer30. *http://www.analog.com/library/analogDialogue/archives/36-03/pll/index.html*, 2002.

[7] Jafar Savoj and Behzad Razavi. A 10-Gb/s CMOS clock and data recovery circuit with a half-rate binary phase/frequency detector. *IEEE Journal of Solid-State Circuits*, 38(1):13–21, 2003.

[8] Borivoje Nikolic, Vojin G Oklobdzija, Vladimir Stojanovic, Wenyan Jia, James Kar-Shing Chiu, and M Ming-Tak Leung. Improved sense-amplifier-based flip-flop: Design and measurements. *IEEE Journal of Solid-State Circuits*, 35(6):876–884, 2000.

[9] Antonio GM Strollo, Davide De Caro, Ettore Napoli, and Nicola Petra. A novel high-speed sense-amplifier-based flip-flop. *IEEE transactions on very large scale integration (VLSI) systems*, 13(11):1266–1274, 2005.

[10] Kenneth K Clarke and Donald T Hess. *Communication circuits: analysis and design*, volume 971. Addison-Wesley Reading, MA, 1971.

[11] Robert B Northrop. *Analog Electronic Circuits: analysis and applications.* Addison Wesley Publishing Company, 1990.

[12] Guillermo Gonzalez. *Foundations of oscillator circuit design.* Artech House, 2007.

[13] Jri Lee, Kenneth S Kundert, and Behzad Razavi. Analysis and modeling of bang-bang clock and data recovery circuits. *IEEE Journal of Solid-State Circuits*, 39(9):1571–1580, 2004.

[14] Youngdon Choi, Deog-Kyoon Jeong, and Wonchan Kim. Jitter transfer analysis of tracked oversampling techniques for multigigabit clock and data recovery. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 50(11):775–783, 2003.

[15] Nicola Da Dalt. Linearized analysis of a digital bang-bang PLL and its validity limits applied to jitter transfer and jitter generation. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 55(11):3663–3675, 2008.

[16] Nicola Da Dalt. Markov chains-based derivation of the phase detector gain in bang-bang PLLs. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 53(11):1195–1199, 2006.

[17] Habib Adrang and Hossein Miar Naimi. Nonlinear Analysis of BBCDR Jitter Generation Using VOLTERRA Series. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 60(4):197–201, 2013.

[18] M Ramezani, C Andre, and T Salama. Jitter analysis of a PLL-based CDR with a bang-bang phase detector. In *Circuits and Systems, 2002. MWSCAS-2002. The 2002 45th Midwest Symposium on*, volume 3, pages III–393. IEEE, 2002.

[19] Jae-Yong Ihm. Stability analysis of bang-bang phase-locked loops for clock and data recovery systems. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 60(1):1–5, 2013.

[20] Habib Adrang and Seyed Saleh Ghoreishi. Modeling of jitter in bang-bang clock and data recovery circuits. *COMPEL-The international journal for computation and mathematics in electrical and electronic engineering*, 32(3):1151–1168, 2013.

[21] Ahmed Gabr and Tad Kwasniewski. Unifying approach for jitter transfer analysis of bang-bang CDR circuits. In *Electronics and Information Engineering (ICEIE), 2010 International Conference On*, volume 2, pages V2–40. IEEE, 2010.

[22] Myeong-Jae Park and Jaeha Kim. Pseudo-linear analysis of bang-bang controlled timing circuits. *Circuits and Systems I: Regular Papers, IEEE Transactions on*, 60(6):1381–1394, 2013.

[23] Richard C Walker. Designing bang-bang PLLs for clock and data recovery in serial data transmission systems, 2003.

[24] Marijn Verbeke, Pieter Rombouts, Arno Vyncke, and Guy Torfs. Influence of jitter on limit cycles in bang-bang clock and data recovery circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 62(6):1463–1471, 2015.

[25] Christophe Van Praet. *Techniques to Reduce Energy Consumption in Next-Generation Access Networks*. PhD thesis, Ghent University, 2014.

5

Experimental results

This chapter presents the results obtained by the measurements performed on the CABINET ASIC. It starts by discussing the experimental test setup, presents the measurements of the chips and concludes with the power consumption estimates from these measurements for CBi-PON. Finally, these estimates are compared according to the GreenTouch reference architecture to show the improvement a CBi-PON can have on the power consumption of Next-Generation communication networks.

5.1 Measurement setup

To verify the correct operation of the CABINET chip, a dedicated test board (Figure 5.1) was developed to provide the required inputs and observe the necessary outputs. The test board was designed such that the CABINET die could be wirebonded directly to the Printed Circuit Board (PCB), eliminating any interconnection parasitics that would degrade the signal quality when using a standard package. The use of a generic test platform developed in the INTEC Design group allows for an easy setup and control from a web browser interface, from which the CABINET can be configured. The connection of the CABINET test board to the generic test platform is shown in Figure 5.2.

The high speed inputs and outputs have been carefully designed using trans-
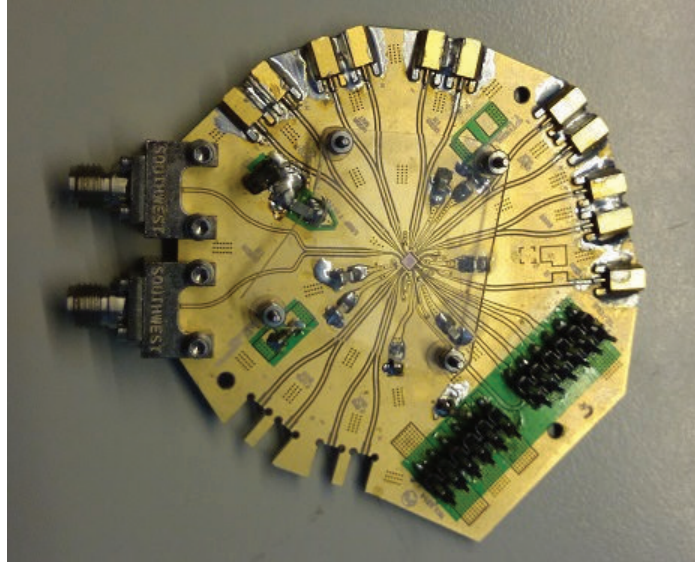
Figure 5.1: CABINET Testboard

mission lines and simulations thereof in Agilent Advanced Design System. Furthermore, the input pair has been connected using 2.92mm Southwest End Launch connectors [1]. These are specified up to 40 GHz, which is sufficient for receiving data streams of 40 Gbit/s. Since the bandwidth of the output pairs is less critical, they use miniSMP connectors [2]. These are characterized by lower performance, but are smaller and less expensive. Finally, a flatcable is used to connect the control signals and the power supplies from the test platform to the CABINET test board.

5.2 Measurement strategy

Due to the immense complexity of the CABINET, not all building blocks can be measured separately: this would require a huge number of test pins, making the chip overly expensive. As a result, measurements must often be done indirectly. For example, the VCO oscillation frequency is not observable directly, yet a divided version is used to clock the serializer. Setting a clock pattern as the input of the serializer allows us to observe a derived version of the desired signal.
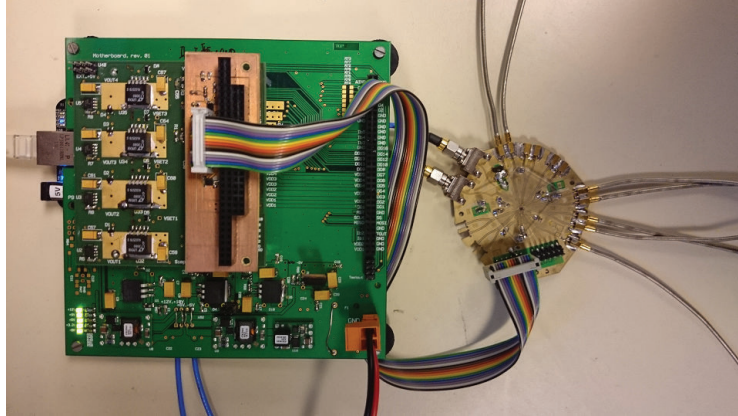
Figure 5.2: CABINET Testboard connected to Testplatform

5.3 Building blocks verification

Before testing the complete CABINET system, it makes sense to start with the verification of the building blocks of the system. This allows us to identify potential issues early in the measurement process and to better understand the behavior of the complete CABINET system.

5.3.1 Voltage-Controlled Oscillator

The VCOs can be measured relatively easy, since the control voltage can be applied externally through the connection for the off-chip loop filter capacitor. The coarse tuning value can be manually set through the I2C control interface.

The output clock can not be observed directly, but a divided version can be measured either on the LVDS clock output of the Medium Access Control (MAC) preprocessor or when serializing a fixed test pattern through the Analog Back-End. Due to some issues with the output serializer described in the next section, all observations were made on the LVDS clock output of the MAC preprocessor, which always results in an output clock 1/8th of the actual oscillation frequency.

For all VCOs, the complete range of coarse values was measured, and for every coarse value the control voltage V_{ctrl} was swept from 0.35 V to 0.85 V in steps of 0.05 V.

The measurement results for the 40 GHz VCO, 10 GHz VCO and 2.5 GHz VCO are shown in Figure 5.3, Figure 5.4 and Figure 5.5 respectively. Both for the 40 GHz and the 10 GHz VCO, not all coarse values resulted in an observable output signal, these measurements were left out of the plots for the sake of clarity.

The unobservable output signal is most likely due to the high oscillation frequency of the VCOs for higher coarse values. However, as the 2.5 GHz VCO can be measured up to the highest coarse value, one can safely assume the 10 GHz VCO is actually oscillating up to a much higher frequency. This assumption stems from the fact that the 10 GHz VCO has the same core as the 2.5 GHz VCO. Therefore, the issue most likely lies in the clock distribution network: while the network and the clock buffers are the same, the frequency of the signal is much higher in case of the 10 GHz VCO which could be problematic for the clock buffers in the tree.

The same reasoning applies for the 40 GHz VCO, which can not be observed at a quarter rate, but is likely to face the same issues as the 10 GHz VCO, as the the oscillation frequency of the VCO is practically the same. This can also be seen from the measurements, as the highest oscillation frequency observed is almost exactly the same in both cases.
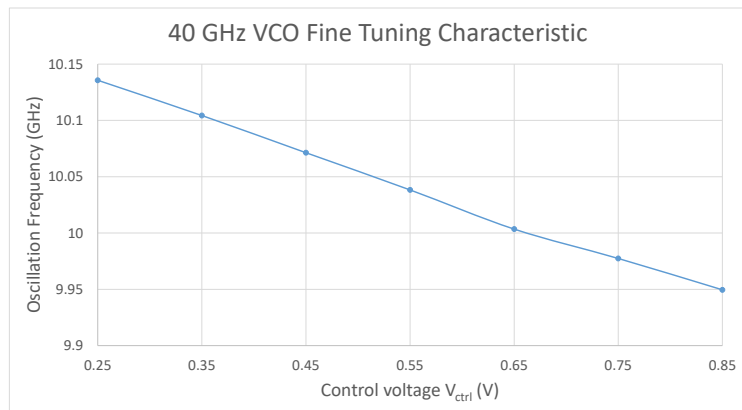
	40 GHz VCO	10 GHz VCO	2.5 GHz VCO
Simulated	-343 MHz/V	-345 MHz/V	-84 MHz/V
Measured	-310 MHz/V	-339 MHz/V	-72 MHz/V

Table 5.1: VCO Gain: Simulated versus Measured at desired oscillation frequency

Nevertheless, all three VCOs clearly demonstrate the ability to oscillate at the required frequency (10 GHz or 2.5 GHz), while also exhibiting the necessary continuous tuning range. Furthermore, the measured VCO gains showed good correspondence with the simulations, as presented in Table 5.1. However, during testing it was observed that the oscillation frequency is in all cases heavily dependent on the supply voltage (e.g. a VCO gain of about 3 GHz/V for the 40 GHz VCO). As a result, relatively limited voltage variations will have a large impact on the oscillation frequency.
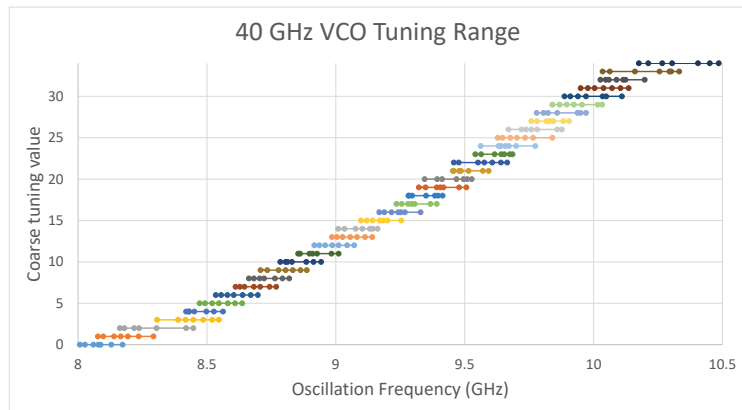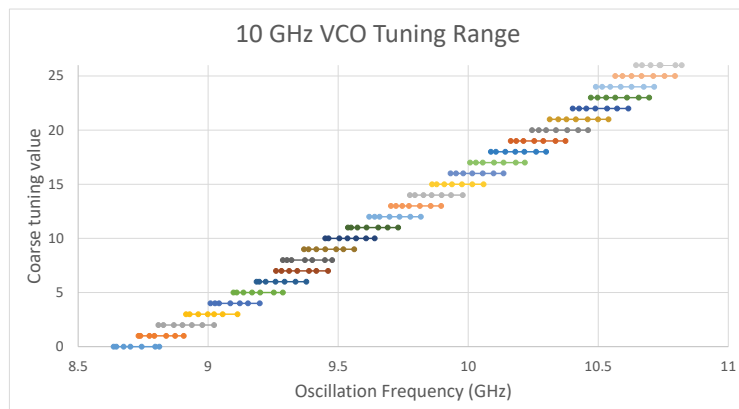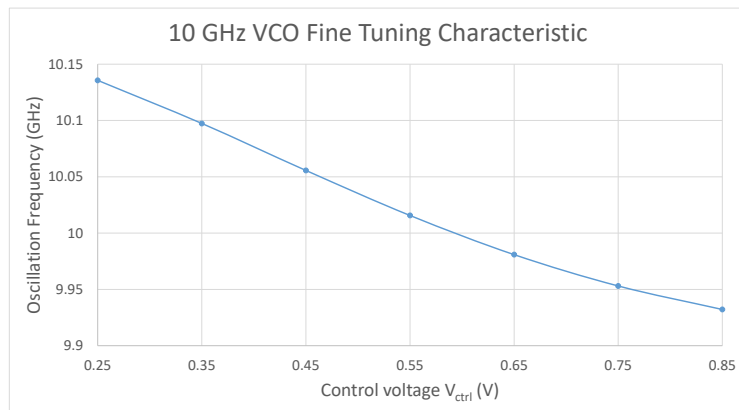
(a) 40 GHz VCO Tuning Range Measurement

(b) 40 GHz VCO Fine Tuning Characteristic
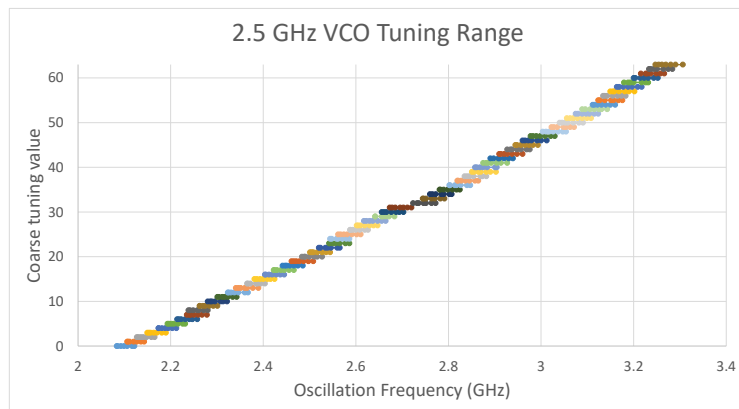
Figure 5.3: 40 GHz VCO Measurement Results

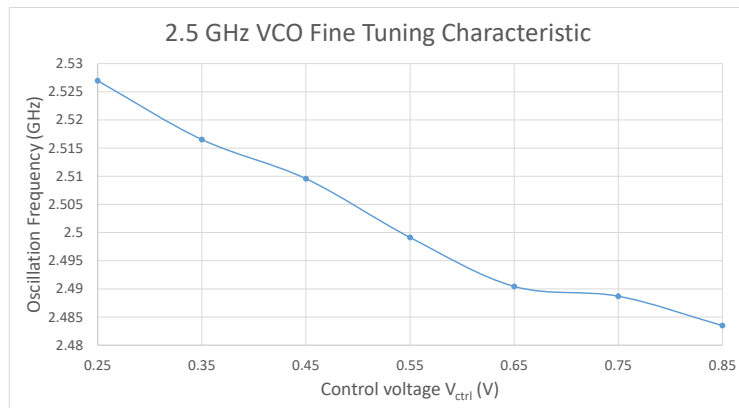(a) 10 GHz VCO Tuning Range Measurement

(b) 10 GHz VCO Fine Tuning Characteristic

Figure 5.4: 10 GHz VCO Measurement Results

(a) 2.5 GHz VCO Tuning Range Measurement

(b) 2.5 GHz VCO Fine Tuning Characteristic

Figure 5.5: 2.5 GHz VCO Measurement Results

5.3.1.1   Output serializer

During testing it was observed that the output serializer, which is part of the Analog Back-End, did not operate as expected. Serializing a fixed test pattern only resulted in an observable, correct output signal when utilizing a 2.5 GHz clock. When serializing with a 10 GHz clock, i.e. with the CABINET configured for 40 Gbit/s or 10 Gbit/s, the serializer did not return any output.

Since the input pattern to the serializer was fixed, no issues are expected there. Furthermore, the serializer does operate correctly for 2.5 Gbit/s, so functionally the design is solid. This means the higher speed of the serializing clock is the root cause of the issue. Extensive simulations were performed on the serializer, trying to reproduce the issue, but none revealed a possible source of the problem. Therefore, we believe the issue comes from the clock distribution network.

In the 10 Gbit/s operating mode, a serializing clock of only 2.5 GHz is expected due to the 1/4 sub-sampling. As 2.5 GHz is a fairly low frequency, the issue in the clock distribution network is expected before the 10 GHz is divided to 2.5 GHz. More extensive simulations of the clock tree confirmed this part of the clock tree to be an area where potential issues could arise.

Unfortunately, the CABINET Repeater Mode was only implemented for 40 Gbit/s and 10 Gbit/s input rates, and not for 2.5 Gbit/s. Therefore, it could never be verified that the serializer works with data supplied from the MAC preprocessor at 2.5 Gbit/s, as this mode was not supported.

The issues with the serializer severely limited the options to test and debug the CABINET. Nevertheless the incorrect operation did provide us with important insights of the failing mechanisms in the CABINET.

5.3.2   Frequency Locked Loop

As the VCOs exhibit the expected behavior, the Frequency Locked Loop (FLL) is the logical next building block to verify. The FLL is the loop that locks the VCO output to a multiple of the reference clock input. The CABINET was configured such that the operation can be tested: when lock is achieved, the CABINET does not automatically switch to the Phase Locked Loop. Furthermore, an external reference clock was applied to the CABINET, since it is the only input signal that is needed.

The control voltage can be observed using the off-chip loop filter pin and should be relatively stable when frequency lock is achieved. Subsequently,

the frequency lock can be verified by viewing the output clock waveform indirectly and measuring its frequency. These measurements have been done for three distinct reference clock frequencies (66 MHz, 72 MHz and 78.125 MHz), showing the FLL locking correctly to the applied reference clock. The base reference frequency 78.125 MHz was chosen such that the FLL locking to the reference frequency results in a 10 GHz or 2.5 GHz recovered clock. The other reference frequencies were chosen in the vicinity of the base reference frequency to demonstrate the capabilities of the FLL.

5.3.2.1   40 Gbit/s FLL

Figure 5.6, Figure 5.7 and Figure 5.8 show the output waveforms of the frequency locked clock with the CABINET configured for 40 Gbit/s operation. For all three reference clock frequencies, locking behavior was observed.
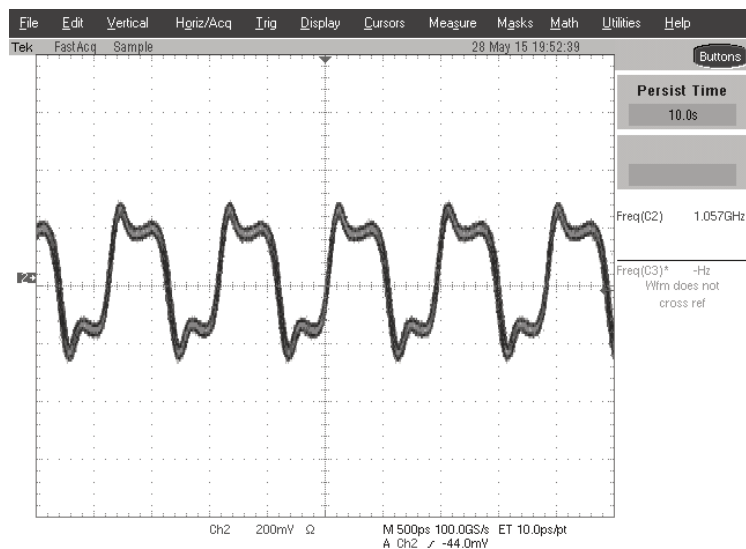
Figure 5.6: Locking to 66 MHz Reference Clock
Expected: : $\frac{66\,\text{MHz} \times 128}{8} = 1.056\,\text{GHz}$
Measured: 1.057 GHz

Figure 5.7: Locking to 72 MHz Reference Clock
Expected: $\frac{72\,\mathrm{MHz}\times128}{8} = 1.152\,\mathrm{GHz}$
Measured: $1.149\,\mathrm{GHz}$

Figure 5.8: Locking to 78.125 MHz Reference Clock
Expected: $\frac{78.125\,\mathrm{MHz}\times128}{8} = 1.250\,\mathrm{GHz}$
Measured: $1.248\,\mathrm{GHz}$

5.3.2.2 10 Gbit/s FLL

The CABINET ASIC configured for 10 Gbit/s input rate shows, like for the 40 Gbit/s case, the expected frequency locking behavior for the external reference clock frequencies 66 MHz, 72 MHz and 78.125 MHz as shown in Figure 5.9, Figure 5.10 and Figure 5.11 respectively.
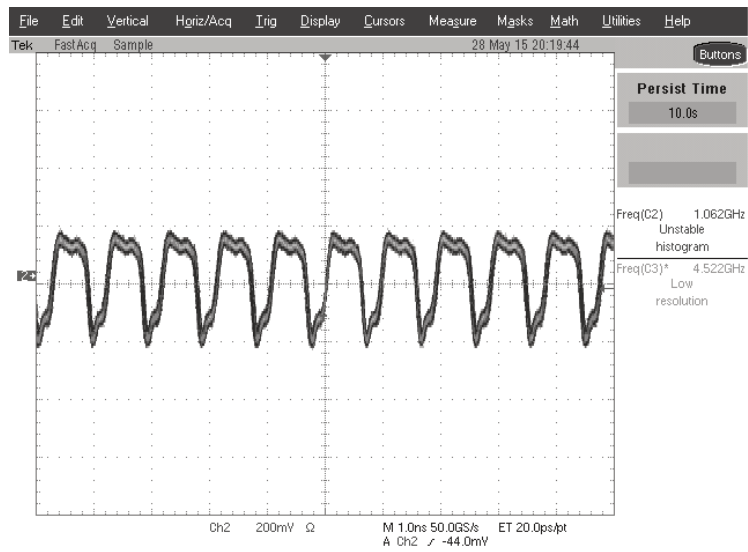
Figure 5.9: Locking to 66 MHz Reference Clock
Expected: $\frac{66\,\mathrm{MHz}\times128}{8} = 1.056\,\mathrm{GHz}$
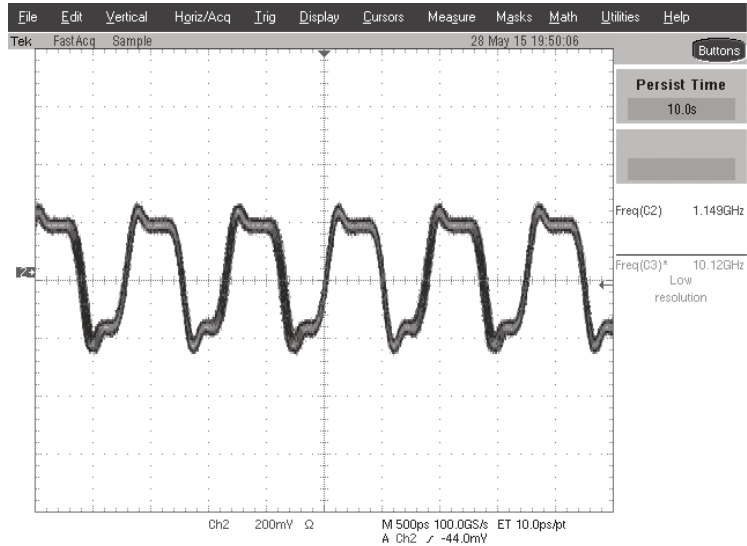Measured: $1.062\,\mathrm{GHz}$
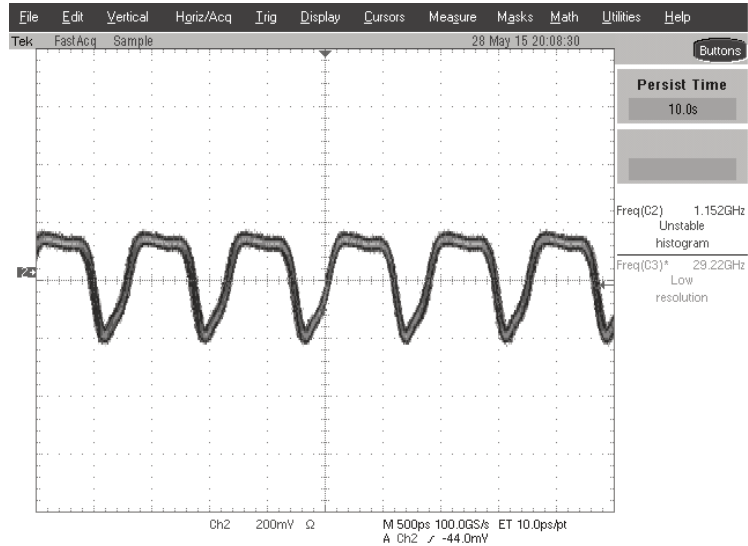
Figure 5.10: Locking to 72 MHz Reference Clock
Expected: $\frac{72\,\text{MHz}\times128}{8} = 1.152\,\text{GHz}$
Measured: $1.152\,\text{GHz}$
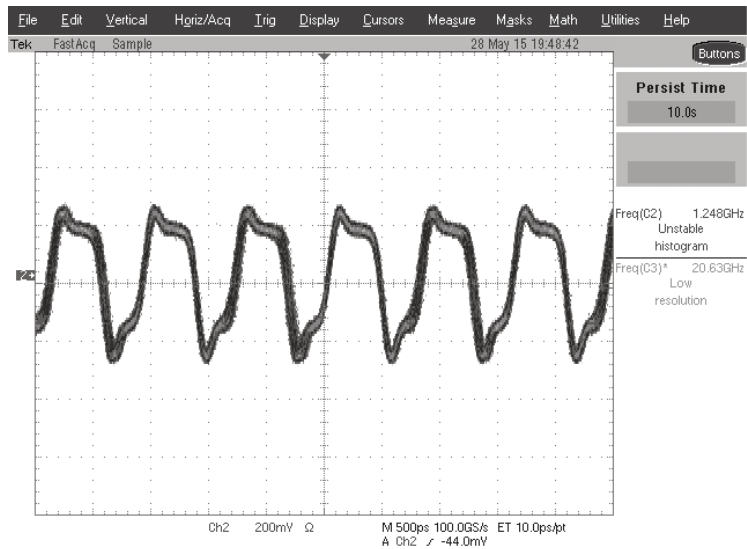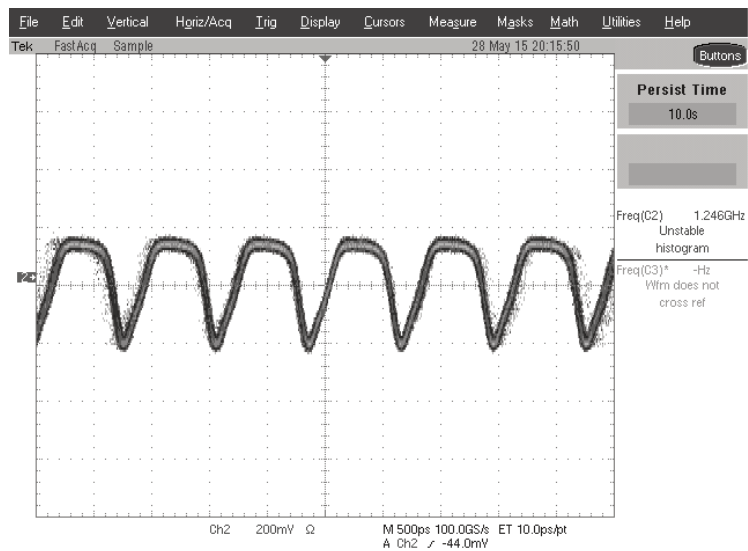

Figure 5.11: Locking to 78.125 MHz Reference Clock
Expected: $\frac{78.125\,\text{MHz}\times128}{8} = 1.250\,\text{GHz}$
Measured: $1.246\,\text{GHz}$

5.3.2.3   2.5 Gbit/s FLL

Finally, the frequency locking experiments confirmed the correction operation of the FLL in the 2.5 Gbit/s configuration, as presented in Figure 5.12, Figure 5.13 and Figure 5.14.
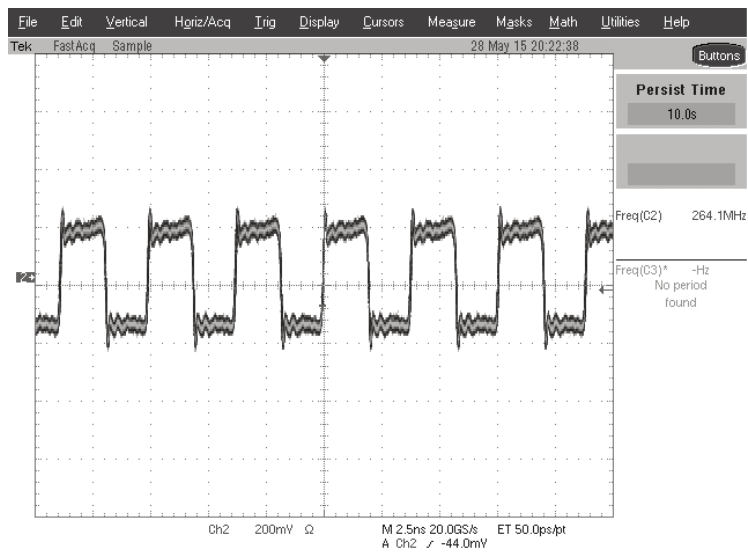
Figure 5.12: Locking to 66 MHz Reference Clock
Expected: $\frac{66\,\text{MHz}\times32}{8} = 264\,\text{MHz}$
Measured: 264.1 MHz
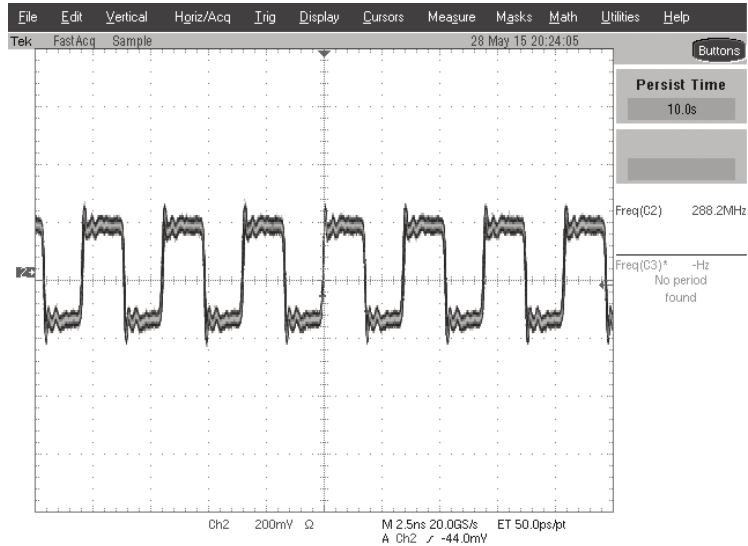
Figure 5.13: Locking to 72 MHz Reference Clock
Expected: $\frac{72\,\mathrm{MHz}\times 32}{8} = 288\,\mathrm{MHz}$
Measured: $288.2\,\mathrm{MHz}$

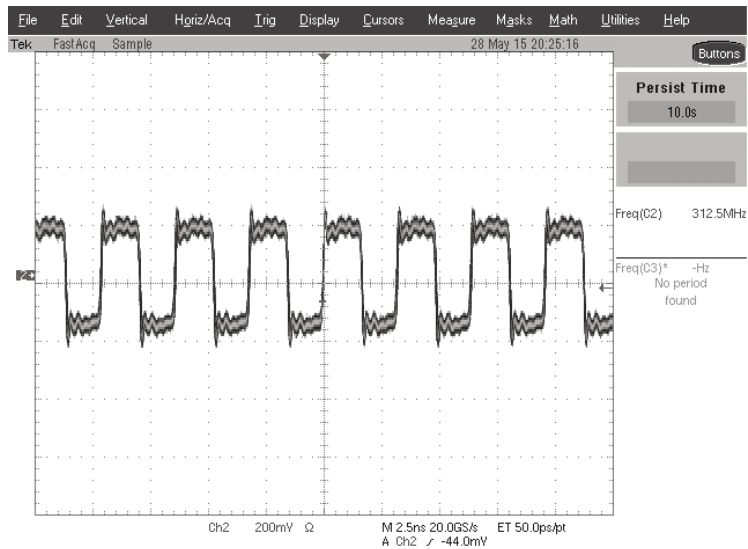
Figure 5.14: Locking to 78.125 MHz Reference Clock
Expected: $\frac{78.125\,\mathrm{MHz}\times 32}{8} = 312.5\,\mathrm{MHz}$
Measured: $312.5\,\mathrm{MHz}$

5.3.3   Samplers

With the FLL verified, the next step is to verify the Clock-and-Data Recovery (CDR). The CDR operates on the input data, so it makes sense to verify that the samplers are working as expected first.

In the CABINET chip, there is no direct path to test the samplers used in the CDR. Of course, applying an input signal is no problem, but verifying that the incoming data is sampled correctly is not as straightforward.

Since the phase detector logic core only runs at 1.25 GHz and is implemented in a hardware description language, it was not identified as being a critical block for the system. Therefore, the CABINET ASIC did not contain means to verify the phase detector logic core.

The sampler operation, however, is considered critical, due to the very high (40 Gbit/s) input data rate. While direct measurement would have been overly expensive in terms of chip area, an indirect path to measure the samplers was implemented. The ASIC was put in this *feedthrough* mode, sending all sampled bits immediately to the output serializer without processing. The serializer output is then observed to see if the bits are correctly recovered.

Due to the VCO sensitivity to the power supply and the serializer speed issues, the samplers could not be measured up to 40 Gbit/s, but the input rate had to be limited to 28 Gbit/s. The high VCO sensitivity to the power supply does influence the measurements, since the shifting VCO phase starts sampling a different channel after some time.

Two patterns were used for testing: `11110000`, which returns the `10101010` pattern regardless of the selected channel. The other pattern used is `0100101100011101`, which results in a different output for every channel, shown in Table 5.2. The first pattern was measured on both 25 Gbit/s and 28 Gbit/s, while the second measurement was only done for 28 Gbit/s.

In Figure 5.15, we see the expected `10101010` output pattern. A faint line is visible at the bottom, presumably when the VCO has shifted too much and a phase change occurs. It is clear the bits are being sampled correctly at this speed.

In Figure 5.16, we again see the expected `10101010` output pattern. A glitch is present, also attributed to VCO phase shifting. This measurement shows that at 28 Gbit/s, the samplers are still able to recover the bits.

Input pattern
0100101100011101

Expected output pattern	
Channel 0	0101010101010101
Channel 1	1001100110011001
Channel 2	0100010001000100
Channel 3	0111011101110111

Table 5.2: Expected output pattern per channel

Using the second pattern, shown in Figure 5.17, the measurement is no longer as clear. Since a VCO phase shift starts sampling a different channel, and the different channels do not output the same pattern, this is to be expected. We see several bits being sampled, but recognizing any pattern is difficult. However, we can clearly identify single bits, which means the samplers are able to sample a single bit without problems.

5.3.4 Clock-and-Data Recovery

Testing of the samplers revealed the sampling of 28 Gbit/s data was not entirely successful, attributed to the low quality of the sampling clock. As a result, measurements on the CDR operation of the CABINET could not produce satisfying results: when the CDR achieves lock, a stable control voltage is expected. However, this was never the case despite numerous attempts to tune the Charge Pump and Loop Filter. Of course, failure to produce a suitable recovered clock also results in faulty recovered data.

5.3.4.1 Switching from FLL to PLL

While Frequency Locking showed the desired behavior, the switch to the Phase Locked Loop introduces some issues. Since the FLL locks to an external reference clock, the input data samplers are not switching at that point. However, when enabling the PLL, the very high speed samplers are toggling high currents. The overall increased power consumption results in a supply voltage drop. As a result the coarse tuning value determined during
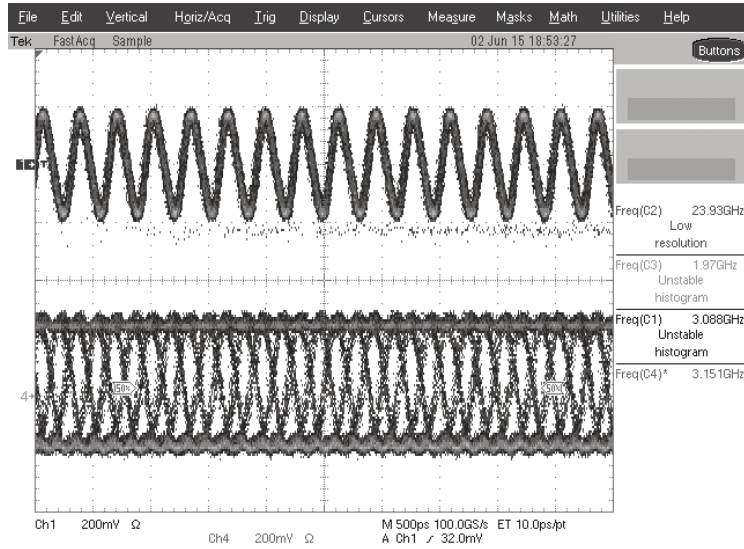
Figure 5.15: 25 Gbit/s Sampler Measurement
Pattern: 11110000
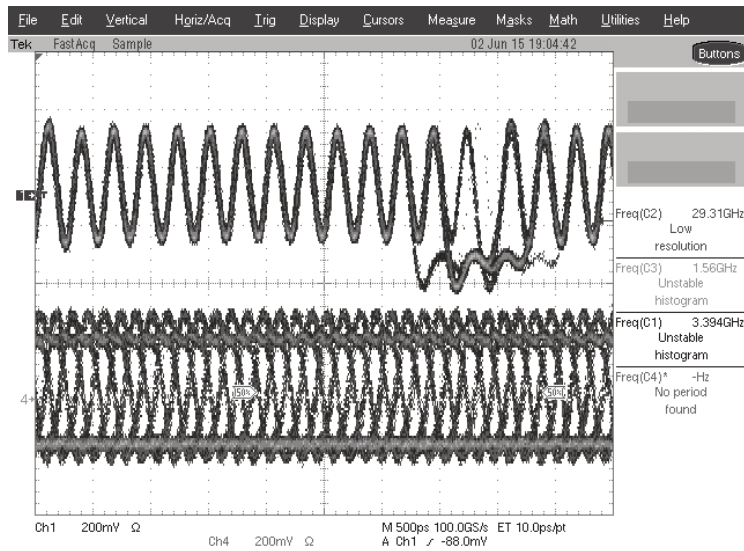
Figure 5.16: 28 Gbit/s Sampler Measurement
Pattern: 11110000
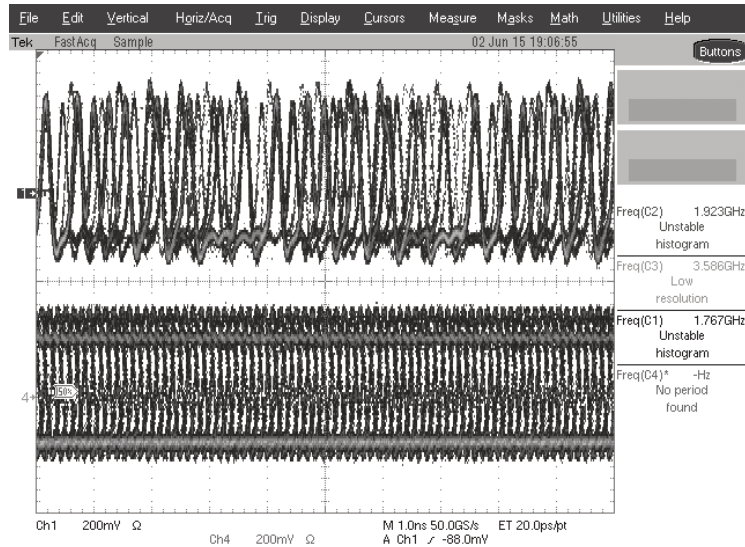
Figure 5.17: 28 Gbit/s Sampler Measurement
Pattern: `0100101100011101`

Frequency Locking is no longer valid when enabling the PLL, which is a first issue with the CABINET CDR. However, by forcing the coarse tuning value manually, this issue can be circumvented.

5.3.4.2 Supply voltage ripple

On top of the voltage drop, the high currents drawn by the toggling high speed samplers impose a significant ripple on the supply voltage. Due to the high power supply sensitivity of the VCO, this ripple directly modulates the oscillation frequency of the VCO, resulting in an enormous increase in jitter. Therefore, the output of the VCO can hardly be regarded a sufficiently clean clock to use for the high speed samplers. As a result, the operation of the CDR is inhibited by the high power supply sensitivity of the VCO.

5.3.4.3 CDR Conclusion

While the simulations showed the CDR loop should operate correctly, these simulations did not take into account the enormous impact of the supply

voltage on the VCO. This prevented the CDR from exhibiting the desired behavior and therefore the CDR could not be used to recover the data.

5.3.5 FLL-based Data Recovery

Even though data recovery was not possible by use of the CDR, the FLL did lock to an external clock frequency correctly. Through trial and error it was determined that 2.5 Gbit/s data streams could be recovered correctly using the clock recovered by the FLL. Of course, operating on 2.5 Gbit/s allows for significantly more jitter on the recovered clock owing to the $4\times$ increase in bit period.

However, the question did arise why the FLL clock could correctly recover the data, but the CDR could not lock. Further investigation revealed the implementation of the PD Logic for the lower rates contains an error, introduced when adjusting the PD Logic from the 40 Gbit/s to the 10/2.5 Gbit/s CDR. Therefore, a locking CDR is out of the question for the 10/2.5 Gbit/s CDR, regardless of the jitter on the recovered clock.

5.4 MAC preprocessor verification

With the help of the FLL-based scheme, the successful recovery of the data at 2.5 Gbit/s means there is one functional mode in which we can test the MAC preprocessor: the 2.5 Gbit/s End-ONT Mode. Unfortunately, there was no Repeater Mode implemented for the 2.5 Gbit/s input rate, and consequently the MAC preprocessor can not be verified in Repeater Mode.

RNID	Payload (hex)	Payload (binary)
38	0xA6	0b10100110
54	0xB6	0b10110110
70	0xC6	0b11000110

Table 5.3: L3 CBi Frame Payloads

To verify the MAC preprocessor, an L3 CBi Frame was constructed containing different payloads for different RNIDs, which was then sent to the

CABINET. The CABINET is configured to operate as a 2.5 Gbit/s End-ONT and to recover the data based on the FLL. The payloads contained in the L3 CBi Frame are listed with their corresponding RNID in Table 5.3.
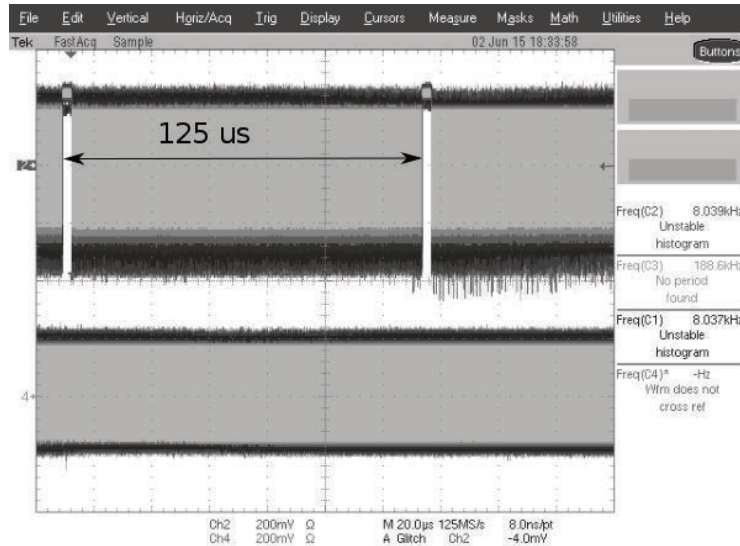
Figure 5.18: 2.5 Gbit/s End-ONT Data Recovery: Frame length

When observing the output of the CABINET, the 125 μs length of the recovered frames was clearly visible, as is shown in Figure 5.18. When observing the recovered payload, it was seen that for each RNID (38, 54 and 70) the expected pattern was recovered as presented in Figure 5.19, Figure 5.20 and Figure 5.21 (in these screenshots the inverted output was observed). As a result, we can conclude the MAC preprocessor operates as expected in the 2.5 Gbit/s End-ONT Mode.
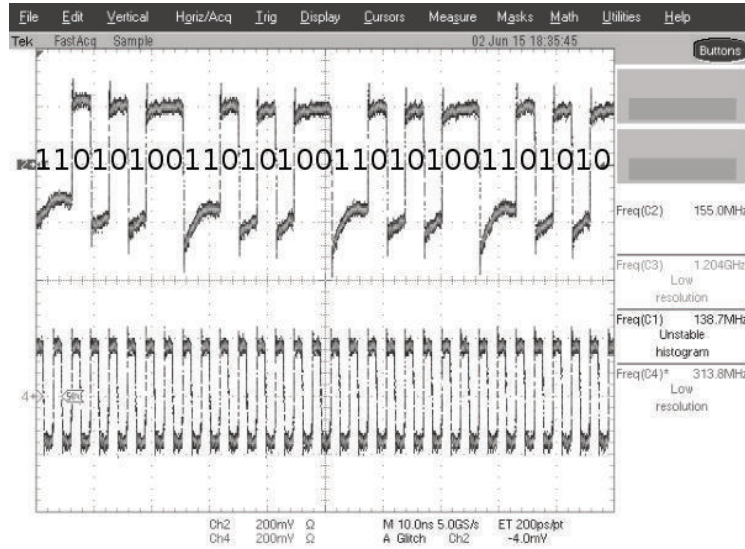
Figure 5.19: FLL-based 2.5 Gbit/s End-ONT Data Recovery
RNID = 38
Payload pattern: `0xA6` = `0b10100110`
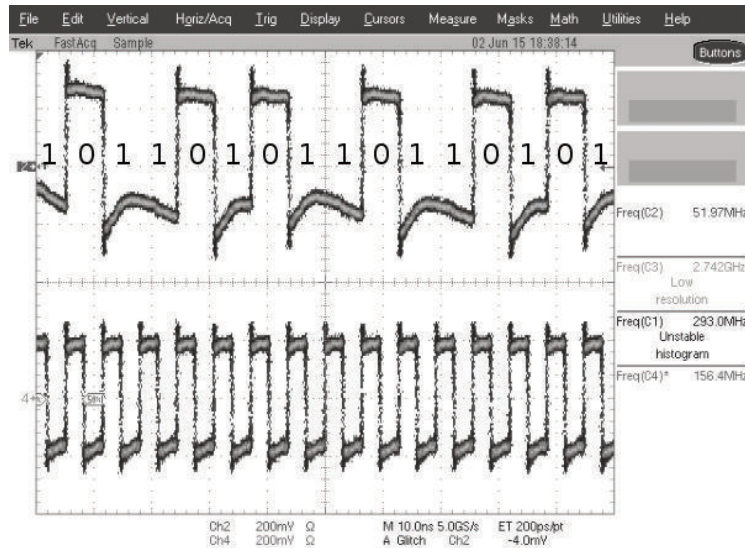
Figure 5.20: FLL-based 2.5 Gbit/s End-ONT Data Recovery
RNID = 54
Payload pattern: `0xB6` = `0b10110110`
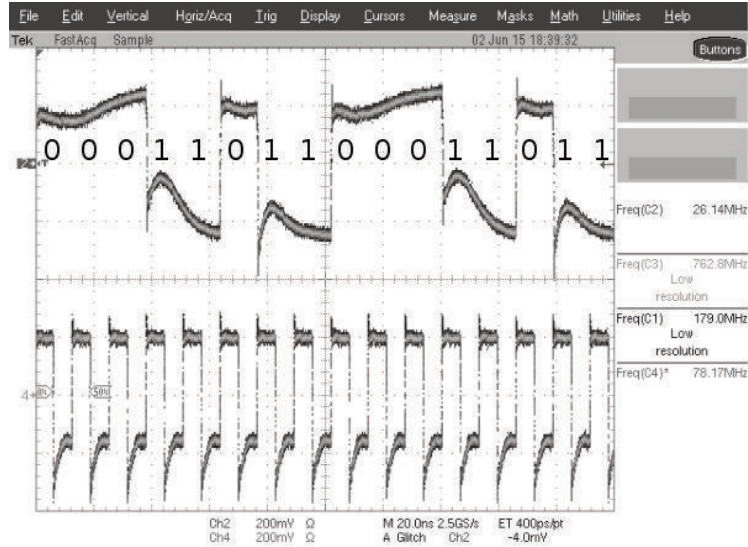
0 0 0 1 1 0 1 1 0 0 0 1 1 0 1 1

Figure 5.21: FLL-based 2.5 Gbit/s End-ONT Data Recovery
RNID = 70
Payload pattern: `0xC6` = `0b11000110`

5.5 Power consumption measurement

Since the 2.5 Gbit/s End-ONT Mode was operating correctly, the power measurements performed on this mode can be used to adjust the power figures extracted from simulation and extrapolate power consumption numbers that are close to reality for a completely functional CABINET.

To do so, the ratio of the simulated power consumption to the measured power consumption is determined. The same ratio is then applied for the other operating modes. Determination of the calibration factors is done separately for each of the power domains. Table 5.4 summarizes the resulting calibration factors.

As noted in Table 5.4, the 2.5 V CML power domain powers only the Analog Back-End and is hence not used in the End-ONT Mode. Therefore, this power domain was measured in 10 Gbit/s Repeater Mode. Despite the incorrect operation of the Analog Back-End, the measured power consumption should be similar to that in correct operation, since the power consumption of the CML blocks does not rely on the switching frequency of the input as is the case for CMOS blocks.

	Analog 1.1 V	Digital 1.1 V	2.5 V LVDS	2.5 V CML[2]
Simulation	20.375 mW	0.5 mW	30 mW	50 mW
Measured	33.96 mW	3.15 mW	24.15 mW	49.225 mW
Calibration	1.7	6.3[1]	0.805	≈ 1

[1] IO buffers could not be simulated, making this a pessimistic calibration factor.

[2] Measured in 10 Gbit/s Repeater Mode.

Table 5.4: Calibration factor determination based on measurements in 2.5 Gbit/s End-ONT Mode

The calibration factors from Table 5.4 are subsequently applied to the power figures coming from the simulations for the other rates and modes of the CABINET, resulting in the extrapolation presented in Table 5.5. Note that the L3 End-ONT does not use any blocks in the 2.5 V CML power domain.

CBi End-ONTs		
	L1 End-ONT	L2 End-ONT	L3 End-ONT
Simulation	106 mW	66.5 mW	50.88 mW
Extrapolated	162.55 mW	94.55 mW	61.94 mW
CBi Repeaters		
	L1 Repeater	L2 Repeater	L3 Repeater
Simulation	126 mW	86.5 mW	not implemented
Extrapolated	187.63 mW	120.48 mW	not implemented

Table 5.5: CABINET power extrapolations

5.6   Power consumption reduction in the network

With those estimates of the power consumption of the CABINET in all modes of operation, it is interesting to investigate the impact on the power consumption of the network. This part of the work has been carried out in the context of the GreenTouch projects CBI and Green Meter Research Study.

In order to do so, the GreenTouch consortium proposed a *2010 baseline network*: a network architecture using the most energy-efficient commercially available technologies at the start of GreenTouch in 2010. This network architecture forms the basis against which the impact of the CBi-PON and the CABINET will be compared.

Note that only residential networks will be investigated here. Business networks typically use a *point-to-point* link from the OLT directly to the ONU. In these types of networks CBi-PON is not relevant.

5.6.1   2010 baseline network (GreenTouch reference network)
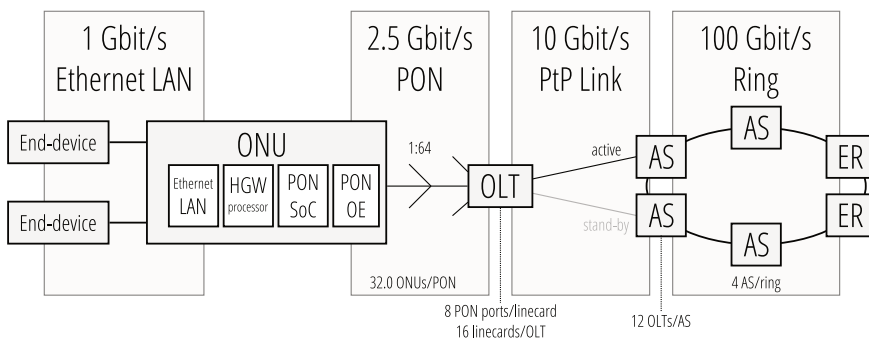
Figure 5.22: 2010 baseline network (GreenTouch reference network)

The 2010 baseline network proposed by the GreenTouch consortium is visualized in Figure 5.22 and serves as the reference network. It identifies two Edge Routers (ER) that connect the metro/aggregation ring to the core network by means of 4 Aggregation Switches (AS). Each AS is then connected directly to 12 OLTs by a point-to-point link. Furthermore, there is a stand-by link to a second Aggregation Switch for redundancy purposes. An OLT consists of 16 linecards with 8 PON ports per linecard. A PON port serves one 2.5 Gbit/s PON with 32 subscribers per PON, which translates to 32 ONUs per PON.

As stated in [3], the power consumption per subscriber in the 2010 baseline network is dominated by the ONU which contributes 6.01 W to the total of 6.726 W or 89.35%. The ONU is clearly the dominant contributor to power consumption per subscriber. As such, it makes sense to focus on the ONU in an effort to improve the energy efficiency of wireline access networks.

5.6.2   2020 GreenTouch network
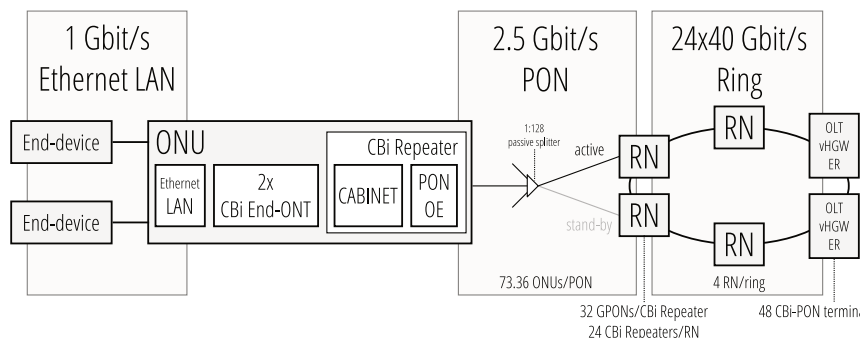
Figure 5.23: 2020 GreenTouch network

Figure 5.23 depicts the residential network architecture that results from all research efforts of the GreenTouch consortium in the area of wireline access networks.

Application of the CBi-PON has a major impact on the network architecture: the extended reach supported by the CBi-PON results in the OLT being co-located with the Edge Router (ER). Furthermore, the Aggregation Switches are replaced by RNs consisting of an array of 24 CBi Repeaters each supporting 32 GPONs.

Moreover, the ONU has changed considerably. The power-hungry SoC used to support GPON operation is now replaced by a low-power 2.5 Gbit/s CBi Repeater which drives the internal 1 Gbit/s PON. This internal PON is terminated by two 1 Gbit/s CBi End-ONTs that effectively extract the traffic for the respective LAN interfaces. This eliminates the need for switching within the ONU.

As a sidenote, the Home Gateway Processor is no longer present in the ONU. One of the GreenTouch technologies that was developed is the *Virtual Home Gateway*, which centralizes the energy-consuming HGW processing and co-locates it with the ER.

5.6.3   Power consumption reduction

Accurately assessing the individual impact of the CABINET on the power savings is very challenging, due to the complexity of the network and the multitude of factors that influence the power consumption. However, in the following sections an effort has been made to clarify the contribution of the CBi-PON.

5.6.3.1   ONU Power consumption

The power consumption of the ONU is directly influenced by the introduction of the CBi-PON: the internal PON consisting of the 2.5 Gbit/s CBi Repeater and the two 1 Gbit/s CBi End-ONTs replaces the power-hungry digital processing SoC.

From [3] we know the SoC consumes 1.481 W in the 2010 baseline network. Calculating the power consumption using the CBi Devices is not as straightforward as it might seem. First, the SoC also supports upstream communication, while the power consumption numbers of the CABINET so far have only dealt with downstream communication. Second, the ONU requires a 2.5 Gbit/s CBi Repeater and a 1 Gbit/s CBi End-ONT. Both were not implemented in the CABINET.

Upstream CBi End-ONTs			
	L1 End-ONT	**L2 End-ONT**	**L3 End-ONT**
Simulation	88 mW	86.5 mW	59.875 mW
Extrapolated	93.93 mW	91.38 mW	60.5125 mW
Upstream CBi Repeaters			
	L1 Repeater	**L2 Repeater**	**L3 Repeater**
Simulation	108 mW	99 mW	not implemented
Extrapolated	128.63 mW	113.325 mW	not implemented

Table 5.6: CABINET Upstream power extrapolations

Luckily, to deal with the first issue, we have simulated power numbers

available for the upstream operation of the CABINET. In short, this implementation consisted of an $8\times$ oversampling CDR due to the burst-mode nature of upstream data traffic. The simulated power numbers were calibrated in the same way as for the downstream case resulting in the power numbers in Table 5.6.

The lack of correct modes of the CABINET to implement the envisioned ONU is mitigated by taking the power numbers for other, higher performance modes: we use the numbers of a 10 Gbit/s (L2) CBi Repeater and those of a 2.5 Gbit/s (L3) CBi End-ONT, which results in a overly pessimistic estimate. Since the power consumption of CMOS circuits is directly related to switching frequency, the power consumption of the CABINET can be scaled according to the operating frequency. Note this means only the 1.1 V power domains are scaled, since the 2.5 V power domain does not use CMOS circuits. The results of this scaling is summarized in Table 5.7.

	L3 End-ONT		L2 Repeater	
Rate	2.5 Gbit/s	1 Gbit/s	10 Gbit/s	2.5 Gbit/s
Power	122.45 mW	93.61 mW	233.805 mW	132.29 mW

Table 5.7: CABINET Power scaling with operating frequency

The numbers reported so far did not take into account the 81% power supply efficiency that GreenTouch uses, nor the $1.5\times$ factor to account for air conditioning. While the power supply efficiency stays the same, the penalty for air cooling is eliminated for the ONU, due to the low power consumption which means passive cooling is sufficient. Equation 5.3 shows the calculation of the ONU power consumption taking into account all these factors, resulting in an ONU power consumption of approximately 0.4 W. Compared to the 1.481 W of the 2010 baseline network this is a power reduction factor of $3.7\times$.

$$P(\text{ONU}) = \frac{P(\text{L2 Repeater})}{0.81} + 2 \times \frac{P(\text{L3 End-ONT})}{0.81} \qquad (5.1)$$

$$\Leftrightarrow P(\text{ONU}) = \frac{132.29\,\text{mW}}{0.81} + 2 \times \frac{93.61\,\text{mW}}{0.81} \qquad (5.2)$$

$$\Leftrightarrow P(\text{ONU}) = 394.46\,\text{mW} \qquad (5.3)$$

5.6.3.2   Remote Node Power consumption

In the 2010 baseline network, Aggregation Switches were used, each consuming about 121 mW/ONU. In the 2020 GreenTouch Network, the aggregation switches have been replaced by Remote Nodes consisting of an array of 32 L1 CBi Repeaters.

The procedure we used to compare these two is very comparable to that for the ONU. The power consumption of the Aggregation Switch includes both downstream and upstream and takes into account the power supply inefficiency. Moreover, the 121 mW/ONU number also includes the power of the optics modules. Therefore, these factors should also be incorporated in the power consumption estimate of the Remote Node that replaces the Aggregation Switch.

Combining the downstream and the upstream numbers, each L1 CBi Repeater consumes 316.26 mW. Each CBi Repeater serves 73.36 subscribers on average. Taking into account the 81% power supply efficiency, this means the CBi Repeater consumes 5.32 mW per ONU. The power consumption numbers for the optics per ONU were obtained through the GreenTouch consortium [3]. The total power consumption per ONU of a Remote Node is summarized in Table 5.8 and amounts to 21.17 mW. Compared to the Aggregation Switch power consumption per ONU in the 2010 baseline network, this translates to a $5.7\times$ power reduction factor.

5.7   Conclusion

In this chapter, the results of the verification of the CABINET ASIC were presented and discussed. The prototype did not operate as expected in all modes: frequency-locking was achieved for all rates, however clock-and-data recovery was not functional. Fortunately, in 2.5 Gbit/s End-ONT mode, data recovery was proven to operate correctly when configured to frequency-lock.

From this 2.5 Gbit/s End-ONT, power consumption numbers could be extrapolated that allowed us to calculate the power consumption reduction compared to the 2010 baseline network defined by GreenTouch. From these calculations, it was shown that the use of CBi End-ONTs reduced the power consumption of the ONU in the 2010 baseline network from 1.5 W to 0.4 W, which is a factor of $3.7\times$.

Moreover, by using the CBi-PON architecture, the power-hungry Aggregation Switches could be replaced by Remote Nodes consisting of an array of

Contributor	Power consumption per ONU
CBi Repeater	$\frac{316.26\,\mathrm{mW}}{73.36\cdot0.81} = 5.32\,\mathrm{mW}$
2.5 Gbit/s OE Optics	4 mW
2.5 Gbit/s OE Electronics	10.16 mW
40 Gbit/s OE Optics	0.57 mW
40 Gbit/s OE Electronics	1.12 mW
Total Remote Node	21.17 mW
Total Aggregation Switch	121 mW
Reduction	5.7×

Table 5.8: Remote Node Power consumption including all factors

CBi Repeaters. As a result, a power reduction factor of 5.7× was achieved at the aggregation point, from 121 mW/ONU to 21.17 mW/ONU.

References

[1] SouthWest Microwave. End Launch Connector Series. `http://mpd.southwestmicrowave.com/products/endLaunch.php`, 2016. [Online].

[2] Rosenberger Hochfrequenztechnik GmbH. Mini-SMP. `http://www.rosenberger.com/us\_en/pdf/products/rf\_coax\_connectors/Mini-SMP.pdf`, 2016. [Online].

[3] GreenTouch Foundation. Improving the energy efficiency of residential fixed access networks by more than $250\times$ by 2020. 2016. [Online].

6

Conclusions and Future work

This dissertation presents the Cascaded Bit-interleaving PON (CBi-PON), a novel PON architecture that supports sustainable growth of next-generation communication networks. The CBi-PON supports high line rates while reducing the power consumption of the communication network significantly. Alongside the fact that a lower power consumption alleviates some of the environmental pressure our global electricity consumption puts on the planet, the operational expenses are also minimized thanks to the lower electricity bill.

Furthermore, the line rate reduces gradually throughout the network owing to the *cascaded* architecture which is implemented by the different CBi Levels. This results in fewer full-rate components needed in the network, which translates to lower capital expenses when installing a CBi-PON when compared to a traditional Long-Reach PON.

In other words, the CBi-PON provides a much-needed solution to deal with the increasing data traffic of the coming years in an energy efficient and cost effective way.

6.1   Summary of this work

In Chapter 1 an introduction was given on the origins of the Internet and its evolution throughout the years. This helps to understand the constraints in

which the work in this dissertation was done.

The Bit-interleaving PON (BiPON) forms the basis on which the Cascaded Bit-interleaving PON (CBi-PON) is built. Since it is such an essential part of the CBi-PON concept, Chapter 2 was used to briefly introduce the reader to BiPON and the crucial paradigm-shift from the traditional packet-based Time-Division Multiplexing (TDM) to the bit-based TDM which is implemented by means of bit-interleaved transmission. The chapter was concluded by the impressive power consumption reduction numbers ($35\times$ up to $180\times$) showing the huge potential of the Bit-interleaving PON in the quest for more energy efficient communication networks.

Chapter 3 started off by highlighting the advantages of a Long-Reach PON as a Next Generation network. Despite the impressive power reductions shown for the Bit-interleaving PON, it was shown that there still is room for improvement by applying the core idea of BiPON to the optics. However, since the optical components that are available today do not support direct sub-sampling, this dissertation has taken the approach of leveraging the mature field of electronics to tackle this power deficiency in BiPON.

Chapter 3 continued to introduce the solution that was developed: the Cascaded Bit-interleaving PON (CBi-PON). The CBi Network Topology was presented (Figure 3.2) and the operation of the CBi Interleaver, CBi Repeater and CBi End-ONT was explained. Subsequently the CBi Frame composition was discussed in detail. Finally, the chapter was concluded with the introduction of a 3-level instantiation of a CBi-PON which is built around a generic CBi Device: the CABINET.

The design and implementation of the CABINET ASIC has been presented in Chapter 4. The system consists of three main parts: the Analog Front-End, the Medium Access Control (MAC) preprocessor and the Analog Back-End. For each of these parts, the design and implementation was discussed in detail, while highlighting the critical building blocks. The chapter was concluded by presenting the final layout of the CABINET.

Finally, Chapter 5 covers the experimental results obtained from the CABINET ASIC for the CBi-PON. The first part of the chapter covers the measurement of the CABINET ASIC and the verification of the different building blocks. Furthermore, it highlights the issues that were encountered and their cause. Despite its issues, it was shown that the CABINET was able to recover the data in the 2.5 Gbit/s End-ONT mode. Using the power consumption of this mode, calibration factors were determined for the relation between the simulated and the measured power numbers. Subsequently, these calibration factors were applied to the simulated power figures of the

non-functional modes of the CABINET to arrive at accurate power estimations for all modes.

With these estimations available, the power consumption reduction in the network could be calculated. To determine the power consumption reduction, the 2010 baseline network defined by GreenTouch was used as the reference network, while the 2020 GreenTouch network was used as the network in which the CBi-PON was deployed. The use of the CBi End-ONTs reduced the power consumption of the ONU in the 2010 baseline network from 1.5 W to 0.4 W, corresponding to a factor of 3.7×. Furthermore, the power-hungry Aggregation Switches in the 2010 baseline network were replaced by Remote Nodes, reducing the power consumption from 121 mW/ONU to 21.17 mW/ONU, or a factor of 5.7×.

6.2 Next-generation networks outlook

The concept of a bit-interleaving PON is clearly a much better choice than the traditional packet-based TDM. Nevertheless, standard committees have not grasped the opportunity presented by NG-PON2, which was originally seen as a disruptive step, to introduce it in the standard. Instead, a conservative network topology was decided upon with the TWDM-PON, which is merely a combination of legacy technologies in order to cope with the increasing bandwidth demands. As a result, it does not address the power issues associated with communication networks.

In doing so, the standard committee has not taken up its responsibility to reduce the tremendous environmental impact of communication networks, but has also made next-generation networks unnecessarily expensive.

6.3 Future work

Regarding the CBi-PON, there still is a lot that can be done. Of course, the current issues with the CABINET should be resolved. This would allow to truly verify the different operating modes of the CABINET and to set up a complete Cascaded Bit-interleaving PON. Additionally, this would yield measured power numbers for all modes of the CABINET, further refining the power consumption reduction figures presented in this dissertation.

Furthermore, it would make sense from a power consumption point of view to design a separate CABINET ASIC for each of the operating modes. In

the current design, building blocks were often designed to support the highest rate, but were also used for much lower rates. While power consumption scales for CMOS circuits, this is not the case for CML circuits. As a result, a significant part of the ONU power consumption can be eliminated by using dedicated building blocks. Due to the complexity of the system and the limited resources, this was not an option for the work that has been presented in this dissertation.

When surpassing the boundaries of the electronics field, one can see that the development of sub-sampling optical receivers would further reduce power consumption and cost in a CBi-PON significantly. Unfortunately, the field of photonics is not mature enough to expect these kinds of optical receivers in the near future.