Evaluatie van de energieconsumptie en het besparingspotentieel in ICT-kernnetwerken

Evaluating the Energy Consumption and the Energy Savings Potential in ICT Backbone Networks

Ward Van Heddeghem

Promotoren: prof. dr. ir. D. Colle, prof. dr. ir. P. Demeester
Proefschrift ingediend tot het behalen van de graad van
Doctor in de Ingenieurswetenschappen: Computerwetenschappen

Vakgroep Informatietechnologie
Voorzitter: prof. dr. ir. D. De Zutter
Faculteit Ingenieurswetenschappen en Architectuur
Academiejaar 2014 - 2015

UNIVERSITEIT
GENT

ISBN 978-90-8578-731-0
NUR 986
Wettelijk depot: D/2014/10.500/77

Universiteit Gent
Faculteit Ingenieurswetenschappen en Architectuur
Vakgroep Informatietechnologie

Promotoren:	prof. dr. ir. Didier Colle prof. dr. ir. Piet Demeester
Juryleden:	prof. dr. ir. Didier Colle, Universiteit Gent (promotor) prof. dr. ir. Patrick De Baets, Universiteit Gent (voorzitter) prof. dr. ir. Piet Demeester, Universiteit Gent (promotor) prof. dr. ir. Lieven Eeckhout, Universiteit Gent prof. dr. ir. Jaafar Elmirghani, University of Leeds dr. ir. Christoph Lange, Deutsche Telekom AG dr. ir. Bart Lannoo, Universiteit Gent prof. dr. ir. Mario Pickavet, Universiteit Gent (secretaris)

Universiteit Gent
Faculteit Ingenieurswetenschappen en Architectuur

Vakgroep Informatietechnologie
Gaston Crommenlaan 8 bus 201, B-9050 Gent, België

Tel.: +32-9-331.49.00
Fax.: +32-9-331.48.99

Proefschrift tot het behalen van de graad van
Doctor in de Ingenieurswetenschappen:
Computerwetenschappen
Academiejaar 2014-2015

Dedicated to Ellemien and Mats,

whom I hope to be able to tell – somewhere in the next years – that the rate at which humankind is consuming our Earth's resources is decreasing instead of increasing...

Dankwoord

Toen ik in 1999 – 15 jaar geleden – afstudeerde als Industrieel Ingenieur Elektromechanica, had ik niet gedacht dat ik nog eens in de academische onderzoekswereld zou belanden. Het had mij nochtans altijd wel wat geleken, 'wetenschapper' worden; ooit gaf ik ergens in het lager onderwijs een tijdje het nogal zelfvoldane antwoord dat ik later 'kernfysicus' wilde worden, want dat klonk erg ingewikkeld en spannend, en het was vermoedelijk ook de job met de moeilijkste naam die ik toen kon bedenken. Maar kijk, uiteindelijk bracht de reis die Het Leven heet mij toch nog langs de wereld van onderzoek, 'papers' en wetenschappelijke conferenties. En alhoewel alles een samenloop is van grote en kleine omstandigheden, wil ik toch graag twee belangrijke momenten vermelden die aan de oorsprong van mijn doctoraatsstudie lagen. Het eerste moment was de beslissing van mijn SGS ex-collega Tom Decreton om na jaren op de arbeidsmarkt toch nog aan een universitaire studie te beginnen; ik volgde prompt zijn voorbeeld en schreef mij in aan de Vrije Universiteit Brussel voor de richting Burgerlijk Ingenieur Toegepaste Computerwetenschappen. Een best spannende beslissing, die ik echter geen seconde heb betreurd. Het tweede moment was in de zomer van 2009, toen ik die studies net afgerond had en voor de zoveelste keer voorbij de lichtjes zweverige maar erg prominente graffiti[1] fietste die jarenlang op de brugpijler aan de Sint-Lievenslaan in Gent prijkte: '*What are you doing to make things better?*' Mijn volgende job kon net zo goed maar beter écht zinvol zijn.

Hier zijn we dan. Vijf jaar later rond ik mijn doctoraal onderzoek rond 'green ICT' af, in een poging om het energieverbruik van o.a. computernetwerken in kaart én omlaag te brengen. Toegegeven, de impact van mijn onderzoek is niet zo wereldschokkend als ik had gedroomd bij de aanvang, maar ik troost mij met de gedachte dat dit vermoedelijk een van de meest gedeelde en herkenbare gemoedstoestanden is onder doctorandi. De 'geïllustreerde gids van een doctoraat' van Matt Might – die het verhaal van zijn zieke zoontje linkt aan de zin van een doctoraalstudie – vat het wat mij betreft mooi samen (zie de figuur hieronder), en hielp mij van tijd tot tijd om de zaken in het juiste perspectief te plaatsen. Misschien dat het sommige van de doctorandi-lezers van dit dankwoord ook kan inspireren;

---

[1] De graffiti is ondertussen verwijderd, maar veelvuldig gefotografeerd en makkelijk terug te vinden op het Internet, bijvoorbeeld op http://goo.gl/4q96Jn.

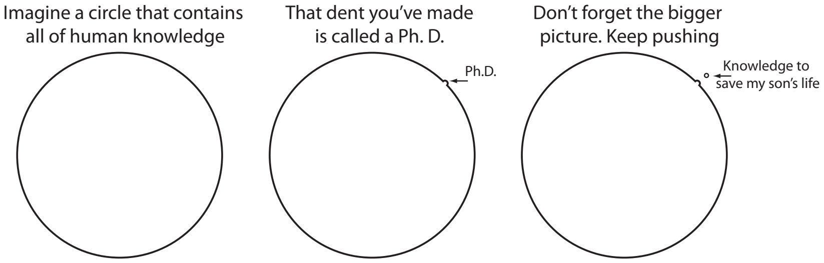

Heel erg beknopte interpretatie van 'The illustrated guide to a Ph.D.' van Matt Might (voor het volledige verhaal, zie http://matt.might.net/articles/phd-school-in-pictures/).

vergeet in dat geval niet de veel uitgebreidere, originele versie te bekijken.

Dit onderzoek was uiteraard niet mogelijk geweest zonder de hulp, het advies en de nodige aanmoedigingen van een aantal personen. Ik wil ze hiervoor dan ook graag bedanken.

Vooreerst wil ik Mario Pickavet bedanken, die dit specifieke onderzoeksonderwerp aanreikte, en mij weer op gang trok toen het allemaal wat vast kwam te zitten. Ik bedank ook met plezier Didier Colle en Piet Demeester die beiden met hun kritische vragen en feedback dit onderzoek van de (soms erg) ruwe kanten ontdeden. Specifiek wil ik Piet trouwens bedanken voor de sfeer van vertrouwen en zelfstandigheid die hij weet te creëren binnen de IBCN onderzoeksgroep.

Ik wil alle mede-auteurs van mijn papers bedanken. In het bijzonder Willem Vereecken, met wie ik het erg fijn vond om samen het model rond de koolstofvoetafdruk van datacenters uit te werken; echt teamwerk komt binnen een doctoraat naar mijn gevoel veel te weinig aan bod, dit was een van de weinige uitzonderingen. Bart Lannoo, die de kunst van het 'deadline flirten' perfect beheerst, maar wel telkens met kritische en nuttige commentaar kwam. I'd like to thank Filip Idzikowski, for the pleasant phone calls we had and his seemingly tireless persistence to try to get things really (really!) right. And Slavisa Aleksic for the nice and interesting collaboration on our Globecom paper. Ik dank ook met veel plezier Sofie Lambert die ik ontelbare keren lastig viel met mijn gemuggenzift over zaken als de meest geschikte lijndikte, tekstkleur of lichtjes alternatieve formulering om een grafiek of tekstfragment toch maar zo duidelijk als mogelijk te krijgen; en waar ze tot mijn genoegen evenveel belang aan hechtte als ikzelf; net als onze gedeelde visie op wetenschappelijke ethiek.

Ik wens ook al mijn (ex-)bureau- en lunchgenoten te bedanken voor de ontspannen babbels en gezellige lunches bij 'de belastingen', in de softmeeting op het 2de verdiep, of in de zon in het park. Het zijn er ondertussen te veel om op te noemen (en om het risico te lopen alsnog iemand te vergeten), maar ze weten wie ze zijn. Bedankt ook Wouter Tavernier voor het delen van o.a. de chaotisch ochtendlijke verhalen over kinderen die, tja, gewoon kind zijn. In the context of my research, I also had the luck and pleasure to work on a couple of international projects, and I would like to thank the various participants from STRONGEST, TREND and GreenTouch for their discussions, the planned and unplanned social events, and the unexpected and slightly adventurous geocaching. Pascal Van Hecke wil ik danken voor de gesprekken op de fiets, niet zelden over 'het werk'; die ritten waren half zo lang, en dubbel zo aangenaam.

Ten slotte kan ik niet anders dan mijn bewondering neerschrijven voor Klaske, mijn vriendin en prachtige moeder van onze twee bengeltjes Mats en Ellemien. Je steun op alle mogelijk momenten, maar zeker bij het naderen van de zoveelste deadline voor een artikel en ik niet altijd op mijn best was, was – en is – onbetaalbaar! Ik kan mij geen betere vriendin wensen!

Gent, september 2014
Ward Van Heddeghem

Table of Contents

Dankwoord iii

Samenvatting xxv

Summary xxix

1 Introduction 1
- 1.1 General introduction to 'Green ICT' . . . 1
  - 1.1.1 The relevance of Green ICT . . . 1
  - 1.1.2 Energy consumption in backbone networks . . . 7
- 1.2 Outline and research contributions . . . 12
- 1.3 Publications . . . 18
  - 1.3.1 Publications in international journals (listed in the Science Citation Index) . . . 18
  - 1.3.2 Publications in other international journals . . . 19
  - 1.3.3 Publications in book chapters . . . 19
  - 1.3.4 Publications in international conferences (listed in the Science Citation Index) . . . 19
  - 1.3.5 Publications in other international conferences . . . 21
  - 1.3.6 Publications in national conferences . . . 22
- References . . . 23

2 Trends in Worldwide ICT Electricity Consumption from 2007 to 2012 27
- 2.1 Introduction . . . 28
- 2.2 Communication networks . . . 30
  - 2.2.1 Telecom operator networks . . . 32
  - 2.2.2 Office networks . . . 34
  - 2.2.3 Customer Premises Access Equipment (CPAE) . . . 35
  - 2.2.4 Results . . . 36
  - 2.2.5 Reliability . . . 37
- 2.3 Personal computers . . . 38
  - 2.3.1 Number of personal computers . . . 39
  - 2.3.2 Number of laptops and desktops, and household and office computers . . . 40

    2.3.3 External monitors 40
    2.3.4 Power consumption per device 41
    2.3.5 Result 42
    2.3.6 Reliability 43
  2.4 Data Centers 43
    2.4.1 Electricity use per server 45
    2.4.2 Worldwide number of servers 45
    2.4.3 Storage, communications and infrastructure overhead 46
    2.4.4 Result 46
    2.4.5 Reliability 47
  2.5 Overall trends and observations 48
    2.5.1 Growth trends 48
    2.5.2 Relative power consumption 51
    2.5.3 Power consumption of the remaining ICT equipment 51
  2.6 Comparison with other studies 52
    2.6.1 Communication networks 54
    2.6.2 Personal computers 55
    2.6.3 Data centers 56
  2.7 Conclusion and outlook 57
    2.7.1 Conclusion 57
    2.7.2 Reflections and outlook beyond 2012 58
  References 60

3 Power Consumption Modeling in Optical Multilayer Networks 65
  3.1 Introduction 66
  3.2 Related Work 68
  3.3 Reference power consumption values 70
    3.3.1 IP/MPLS layer 71
    3.3.2 Ethernet Layer 74
    3.3.3 OTN layer 75
    3.3.4 WDM layer 76
  3.4 Analytical power consumption model 77
    3.4.1 Model 79
    3.4.2 Explanation 81
    3.4.3 Comparing with earlier analytical models 83
  3.5 Evaluation and case study 84
    3.5.1 Cases considered 84
      3.5.1.1 Network topologies 84
      3.5.1.2 Network traffic demands 84
      3.5.1.3 Node architecture and (non-)bypass scenarios 86
      3.5.1.4 Calculation methods 88
    3.5.2 Results 91
      3.5.2.1 Model evaluation 91
      3.5.2.2 Component power consumption distribution 92
      3.5.2.3 Savings from optical bypass 95

3.6 Conclusions . . . 97
References . . . 99

4 A Power Consumption Sensitivity Analysis of Circuit-Switched Versus Packet-Switched Backbone Networks 103

4.1 Introduction . . . 104
4.2 Related work . . . 107
4.3 Network architecture: circuit switching versus packet switching . . . 108
4.4 Network dimensioning and power consumption calculation 110
  4.4.1 Power consumption model . . . 110
  4.4.2 Dimensioning and power consumption calculation . 111
4.5 Case-studies - network scenarios . . . 113
  4.5.1 Topologies . . . 113
  4.5.2 Traffic matrices . . . 115
  4.5.3 Linerates . . . 115
4.6 Results and observations . . . 115
  4.6.1 General observations . . . 116
  4.6.2 Influence of the demand/linerate ratio . . . 118
  4.6.3 Influence of the network size (number of nodes and physical link lengths) . . . 120
  4.6.4 Influence of the mesh degree . . . 123
  4.6.5 Sensitivity to non-uniform demands . . . 125
  4.6.6 Sensitivity to a more detailed IP power consumption calculation . . . 126
4.7 Conclusion and further work . . . 128
References . . . 131

5 A Quantitative Survey of the Power Saving Potential in IP-over-WDM Backbone Networks 135

5.1 Introduction . . . 136
5.2 Related work . . . 141
5.3 Methodology . . . 143
  5.3.1 Baseline scenario and reference power rating values . 144
  5.3.2 Analytical power model . . . 146
  5.3.3 Calculating power savings . . . 148
    5.3.3.1 Savings for one equipment type . . . 148
    5.3.3.2 Savings across multiple equipment types . . 148
  5.3.4 Expressing power savings . . . 149
  5.3.5 The Moderate Effort and Best Effort scenario . . . 150
5.4 Approaches to save power . . . 150
  5.4.1 External overhead factor η_{eo} . . . 152
  5.4.2 Protection factor η_{pr} . . . 156
  5.4.3 Amount of traffic T . . . 160
  5.4.4 Power rating P/C (part I) . . . 167

5.4.5 Power rating P/C (part II) . . . 171
5.4.6 Hop count H . . . 179
5.4.7 Overprovisioning factor η_{op} . . . 181
5.5 Total power saving potential and discussion . . . 182
5.5.1 Total power saving potential . . . 183
5.5.1.1 Static approaches only . . . 183
5.5.1.2 Static and time-dependent approaches . . . 185
5.5.2 Sensitivity to the average link length . . . 186
5.5.3 Sensitivity to the hop count H . . . 186
5.5.4 Comparison with Green Meter . . . 188
5.5.5 The role of legacy networks . . . 190
5.6 Conclusion and recommendations . . . 191
References . . . 194

6 Conclusion **207**

A Distributed Computing for Footprint Reduction by Exploiting Low-Footprint Energy Availability **213**
A.1 Introduction . . . 214
A.2 Related work . . . 217
A.3 Theoretical model . . . 218
A.3.1 Usage footprint . . . 220
A.3.2 Manufacturing footprint . . . 221
A.3.3 Communication footprint . . . 222
A.3.4 Total footprint . . . 223
A.4 Parameter quantification . . . 224
A.4.1 Manufacturing fraction (f) . . . 224
A.4.2 High-footprint energy emission intensity (I_H) . . . 225
A.4.3 Low-footprint energy emission intensity (I_L) . . . 225
A.4.4 Low-footprint energy availability (p) . . . 226
A.5 Case study I: The Added Distributed Data Centers (ADD) Scenario . . . 227
A.5.1 Influence of manufacturing fraction (f) . . . 227
A.5.2 Influence of low-footprint energy availability (p) . . . 228
A.5.3 Influence of n and m values . . . 230
A.5.4 Influence of emission intensity difference . . . 231
A.6 Case study II: The Low Load Redistribution (LLR) Scenario . . . 233
A.7 Conclusions . . . 234
References . . . 238

List of Figures

1.1 The worldwide electricity consumption of Information and Communication Technology (ICT) during operation (Pickavet, 2008) 4
1.2 Evolution of domestic prices of electricity (incl. taxes) in the EU from 1995 to 2011 5
1.3 The processor power density has been increasing exponentially 6
1.4 Trends in the heat density in computer systems and telecommunications equipment 6
1.5 Two examples of environmentally driven major initiatives to reduce the power consumption of ICT equipment 8
1.6 The number of publications on energy-efficient wired networks has increased significantly in the last years 8
1.7 Evolution of the national percentage of individuals using the Internet 9
1.8 High-level and illustrative breakdown of the architecture of the Internet into backbone networks, metro networks and access networks 10
1.9 Evolution of the architecture of backbone networks 11
1.10 Power consumption per customer in the access and backbone network 13
1.11 Schematic overview of the main publications and their relation to the five main chapters in this dissertation 14

2.1 The scope of ICT equipment that we consider in this study. . 31
2.2 Worldwide use phase electricity consumption of communication networks 37
2.3 Worldwide use phase electricity consumption of personal computers 42
2.4 Worldwide use phase electricity consumption of data centers 47
2.5 Evolution of worldwide electricity use of networks, PCs and data centers and total worldwide electricity use 49
2.6 A comparison of our 2012 results with a number of other often-cited studies 53

3.1 Simplified IP/MPLS router block model 72

3.2 Core router power distribution among the different components . . . 72
3.3 Core router power consumption as a function of the total node capacity, for maximally equipped configurations 74
3.4 Required router ports for one 1+1 protected demand 81
3.5 Required optical line amplifiers for one 1+1 protected demand 83
3.6 IP topologies of the test networks 85
3.7 Network node and link architecture 87
3.8 Power consumption with increasing traffic demands (gravity, random and uniform) for the pan-EU and NSFNET network. 93
3.9 Comparison of the component power consumptions as calculated by the simulation approach, and the power model approach 94
3.10 Detail of inset in Fig. 3.8 94
3.11 Relative component power consumption for the pan-EU network at 80 Gbps average demand (simulation results) 95
3.12 Component power consumption for the pan-EU network (gravity matrix, 80 Gbps) 95
3.13 Relative savings of bypass over non-bypass (pan-EU network, gravity matrix) 96

4.1 The packet-switched and circuit-switched network architectures considered in this paper 108
4.2 The total power consumption and inverse power efficiency of a 15-node ring and full-mesh topology with increasing node-to-node traffic demand 117
4.3 Power savings of CS over PS mapped to the ratio of the demand bitrate over the channel linerate 118
4.4 Influence of the node count on the power savings of CS over PS 121
4.5 Influence of the average physical link length on the relative savings of CS over PS 122
4.6 Influence of the mesh degree on the relative savings of CS over PS 123
4.7 Influence of the mesh degree and the demand/linerate ratio d/l on the relative savings of CS over PS 124
4.8 Influence of different demand types on the savings of CS over PS, both for the (a) Géant and (b) Abilene topology . . . 126
4.9 Influence of using more accurate IP power consumption modelling on the total power consumption 127
4.10 Influence of using more accurate IP power consumption modelling on the power savings of CS over PS 128

5.1 Traffic growth and new equipment energy efficiency improvement in the IP backbone, normalized to the year 2010 137

5.2 Conceptual model of our generalized IP-over-WDM network architecture . . . 145
5.3 Overview of power reduction approaches mapped to the general form of our analytical model . . . 151
5.4 Rough distribution of the external overhead power among its main contributors . . . 153
5.5 Normalized (to baseline) power consumption of several power saving algorithms as reported for the COST239 network . . 158
5.6 Power scaling with load and capacity . . . 168
5.7 Evolution of power rating values of various telecom equipment as given by Neilson, Han, Tamm, and Lange . . . 170
5.8 Theoretical achievable sleep mode savings S for a sinusoidal traffic variation over the day . . . 175
5.9 Reduction factors for the Moderate Effort and Best Effort scenarios for (a) the static-approaches only, and (b) the static and time-dependent approaches . . . 184
5.10 Sensitivity of the combined reduction factor to the average link length . . . 187
5.11 Sensitivity of the combined reduction factor to the average hop count H . . . 187

A.1 Distributed data center . . . 216
A.2 Distributed data center infrastructure overview, consisting of m=5 sites with n=3 sites active. The independent LF energy availability per site is p=0.6 . . . 219
A.3 Lower and upper emission intensity estimates for various low-footprint sources [9] (CSP: Concentrated Solar Power) . 226
A.4 The total normalized footprint F_{norm} and corresponding relative emission savings as a function of the total number of data centers m . . . 228
A.5 The total normalized footprint F_{norm} and relative emission savings for n=3 as a function of m for different manufacturing fractions f . . . 229
A.6 The total normalized footprint and the relative savings (with respect to the baseline scenario where m=n=3) as a function of the LF energy availability p (n=3, m=6 and f=0.25) 230
A.7 Relative emission savings for n=10 (all other parameters are equal as before) . . . 231
A.8 Relative emission savings as a function of the HF energy emission intensity I_H. (for f=0.25, p=0.6 and I_L=10 g CO_2-eq/kWh) . . . 232
A.9 Relative emission savings as a function of the LF energy emission intensity I_L. (for f=0.25, p=0.6 and I_H=500 g CO_2-eq/kWh) . . . 232
A.10 Nominal load distribution in a distributed data center 234

A.11 Relative footprint of a distributed data center running at various loads both under a regular scenario and a Low Load Redistribution (LLR) scenario 235
A.12 Savings for various values of p 235

List of Tables

2.1 Office networks: cooling overhead factors and worldwide electricity use per type of equipment 35
2.2 Personal computers and computer monitors: average energy consumption per device and worldwide electricity use per type of equipment 39
2.3 Worldwide power consumption of data centers in 2007 and 2012 44
2.4 Comparison of historical and projected annual growth rates in worldwide electricity use 54
2.5 Comparison of the average energy consumption (kWh/year) per PC (inc. monitor) 56

3.1 IP/MPLS components 73
3.2 Ethernet layer (bidirectional) 75
3.3 OTN layer (bidirectional) 75
3.4 WDM components (bidirectional) 78
3.5 Symbols and values 80
3.6 Network topology parameters 85
3.7 Traffic matrices 86
3.8 Network architectural parameters 88
3.9 Dimensioning via simulation power values 89
3.10 Analytical power model valuess 89

4.1 Power consumption values 110
4.2 Notation used in the network dimensioning 112
4.3 Topologies considered in this study 114
4.4 Overview of our findings. Key findings in bold 116
4.5 Calculating the mesh degree for the 3 networks considered in [6] confirms the finding that higher mesh degrees correspond to reduced preference for CS 124

5.1 An overview of surveys on energy efficiency in (backbone) telecommunication networks 138
5.2 Power rating values and weighting factors, reference year 2010 146
5.3 Conversion table for the reader's convenience 150

5.4 Indicative national operator PUE values for determining TNI PUE values 155
5.5 Reduction factor for external overhead 156
5.6 Protection factor related energy savings reported in publications 157
5.7 Reduction factor for protection 161
5.8 Traffic factor related energy savings reported in publications 162
5.9 Reduction factor for amount of traffic 166
5.10 Power rating factor related energy savings reported in publications (part I) 167
5.11 Power rating factor related energy savings reported in publications (part II) 172
5.12 Reduction factor for the power rating factor 179
5.13 Hop count related energy savings reported in publications . 180
5.14 Reduction factor for hop count 181
5.15 Overprovisioning factors reported in publications 182
5.16 Reduction factor for overprovisioning 183
5.17 GreenTouch's Green Meter power reduction estimations . . 188

A.1 Manufacturing fraction values according to different studies 222
A.2 Average CO_2 emissions per kWh from electricity and heat generation for a number of countries and regions 225

List of Acronyms

A

ADD	Added Distributed Data centers
ADM	Add/Drop Multiplexer
ATM	Asynchronous Transfer Mode

C

CAGR	Compound Annual Growth Rate
CapEx	Capital Expenditure
CCN	Content-centric Networking
CDN	Content Distribution Network
CFP	C Form-factor Pluggable
CMOS	Complementary Metal-Oxide Semiconductor
CO	Central Office
CPAE	Customer Premises Access Equipment
CPE	Customer Premises Equipment
CRT	Cathode Ray Tube
CS	Circuit Switching
CSP	Concentrated Solar Power

D

DCF	Dispersion-compensating Fiber
DC	Direct Current

DC	Data Center
DGE	Dynamic Gain Equalizer
DPP	Dedicated Path Protection
DSL	Digital Subscriber Line
DSP	Digital Signal Processing
DWDM	Dense Wavelength Division Multiplexing

E

EA	Energy Aware
EAAR	Energy Aware Adaptive Routing
EA-DPP	Energy Aware Dedicated Path Protection
EAR	Energy Aware Routing
EASP	Energy-Aware Shared Backup Protection
EASPP	Energy-Aware Shared Path Protection
ECR	Energy Consumption Rating
EDFA	Erbium Doped Fiber Amplifier
EEE	Energy-Efficient Ethernet
EON	Elastic Optical Network
EPA	Environmental Protection Agency
EPAR	Energy Profile Aware Routing
EU	European Union
EUSPP	Energy Unaware Shared Path Protection
EWA	Energy Watermark Algorithm

F

FCS	Fabric Card Shelf
FT	France Telecom
FTSFTW	Follow The Sun/Follow The Wind
FTTH	Fiber To The Home
FUFL	Fixed Upper Fixed Lower

G

GA	Genetic Algorithm
Gbps	Gigabit per second
GHG	Greenhouse Gas
GPON	Gigabit Passive Optical Network
GPS	Global Positioning System

H

HDTV	High-definition Television
HES	High-end server
HEVC	High Efficiency Video Coding
HF	High-Footprint

I

IBBT	Interdisciplinary Institute for Broadband Technology
IBCN	Internet Based Communication Networks and Services
ICT	Information and Communication Technology
IDC	International Data Corporation
IEA	International Energy Agency
IEEE	Institute of Electrical and Electronics Engineers
ILP	Integer Linear Programming
IP/MPLS	Internet Protocol/Multiprotocol Label Switching
IP	Internet Protocol
IPCC	Intergovernmental Panel on Climate Change
ISP	Internet Service Provider
ITU	International Telecommunication Union

L

LAN	Local Area Network

LCA	Life Cycle Analysis
LCD	Liquid Crystal Display
LCS	Line Card Shelf
LC	Line Card
LF	Low-Footprint
LFA	Least Flow Algorithm
LLR	Low Load Redistribution
LP	Linear Programming

M

MEMS	Microelectromechanical Systems
MILP	Mixed-Integer Linear Programming
MPLS	Multiprotocol Label Switching
MRS	Mid-range server
MSPP	Multiservice Provisioning Platform
MUELL	Maximum Utilization of Each Logical Link

O

OADM	Optical Add/Drop Multiplexer
OECD	Organisation for Economic Co-operation and Development
OEO	Optical-Electrical-Optical
OLA	Optical Line Amplifier
OTN	Optical Transport Networking
OXC	Optical Cross-Connect

P

PAR	Power Aware Routing
PASPP	Power-Aware Shared Path Protection
PC	Personal Computer
PDA	Personal Digital Assistant

PDU	Power Distribution Unit
PON	Passive Optical Network
POTP	Packet Optical Transport Platform
PoS	Packet-over-SONET/SDH
PS	Packet Switching
PUE	Power Usage Effectiveness

Q

QoP	Quality of Protection
QoS	Quality of Service

R

ROADM	Reconfigurable Optical Add-Drop Multiplexer

S

SBP	Shared Backup Protection
SDH	Synchronous Digital Hierarchy
SDTV	Standard-definition Television
SFP	Small Form-factor Pluggable
SLA	Service Level Agreement
SONET	Synchronous Optical Networking
SPP	Shared Path Protection
SP	Shortest Path
SP-DPP	Shortest Path Dedicated Path Protection
STM	Synchronous Transport Module

T

TID	Telefónica I+D

TM	Traffic Matrix
TNI	Telecom Network Infrastructure
TREND	Towards Real Energy-efficient Network Design
TXP	Transponder

U

UHD	Ultra High Definition
UN	United Nations
UPS	Uninterruptible Power Supply

V

VCR	Videocassette Recorder
VS	Volume server

W

WDM	Wavelength Division Multiplexing
Wi-Fi	Wireless Fidelity
WLAN	Wireless Local Access Network

X

XC	Cross-Connect
XFP	10 Gigabit Small Form-factor Pluggable

Samenvatting
– Summary in Dutch –

De voorbije decennia is de (wetenschappelijke) consensus ontstaan dat de invloed van de mensheid op ons natuurlijk ecosysteem gegroeid is tot voorbij haar draagkracht. Onze ecologische voetafdruk wordt nu geschat op 1,5x de aarde, wat betekent dat onze planeet nu een jaar en zes maanden nodig heeft om te regenereren wat we hebben verbruikt in één jaar. Het is duidelijk dat we de reserves aanspreken van onze natuurlijk bronnen—zoals water, hout, propere lucht, of energie—die opgebouwd zijn gedurende de voorbije honderden of duizenden jaren, of zelfs nog langer wanneer het over olie en gas gaat. En hoewel het ongebreidelde gebruik van elk van deze natuurlijk bronnen onze aandacht verdient, ligt de focus van dit werk op de elektriciteitsconsumptie van Informatie en Communicatie Technologie (ICT). ICT is een brede term om elektronische goederen en diensten aan te duiden die gerelateerd zijn aan gegevensverwerking en informatieoverdracht.

In het eerste deel van dit werk schatten we het wereldwijde elektriciteitsverbruik van ICT toestellen in. Aansluitend bekijken we hoe het elektriciteitsverbruik in ICT kernnetwerken kan worden gereduceerd. Tenslotte evalueren we het gebruik van zonne-energie en windenergie om de uitstoot van broeikasgassen in gedistribueerde datacenters te beperken.

Inschatten van het elektriciteitsverbruik van ICT Het uitgangspunt van deze thesis was de observatie in een eerdere studie dat het elektriciteitsverbruik van ICT toestellen jaar na jaar toenam. Voor 2008 werd het gecombineerde groeitempo van het verbruik van communicatienetwerken, personal computers en datacenters geschat op 10% per jaar (een verdubbeling elke 9 jaar), wat sneller is dan de 3% per jaar waarmee het wereldwijde algemene elektriciteitsverbruik groeit. Gezien diezelfde drie categorieën in datzelfde jaar al 4% van het wereldwijde algemene elektriciteitsverbruik voor hun rekening namen, zou dit betekenen dat ze in 2040 evenveel elektriciteit zou verbruiken als de helft van het wereldwijde algemene elektriciteitsverbruik in 2008. Het is duidelijk dat deze toename vanuit een ecologisch oogpunt noch houdbaar, noch duurzaam is.

Een eerste bijdrage van deze thesis is het inschatten hoe de jaarlijkse groei in het ICT elektriciteitsverbruik veranderd is sinds de studie uit 2008; vooral gezien het feit dat er de voorbije jaren een toegenomen aandacht is voor het verbeteren van de energie-efficiëntie in verschillende sectoren, waaronder ook ICT. Hiervoor bepalen we het elektriciteitsverbruik van communicatienetwerken, personal computers en datacenters in 2007 en 2012. Onze resultaten tonen aan dat de gecombineerde groei van deze drie categorieën nu 7% per jaar bedraagt. Dit is lager dan de 10% jaarlijkse groei die vóór 2007 werd vastgesteld voor dezelfde drie categorieën. Het elektriciteitsverbruik van ICT groeit dus inderdaad minder snel dan het geval was vóór het jaar 2007. Een belangrijke reden hiervoor is een verschuiving naar meer energie-efficiënte technologieën, zoals van desktop computers naar laptops, van omvangrijke CRT monitors naar platte LCD schermen, en de doorbraak van server virtualisatie en efficiëntere koeling in datacenters. Hoewel dit goed nieuws is, mogen we niet uit het oog verliezen dat het ICT elektriciteitsverbruik nog steeds een pak sneller groeit dan het wereldwijde algemene elektriciteitsverbruik. Elk van de drie categorieën—communicatienetwerken, personal computers, en datacenters—verbruikt grofweg een zelfde hoeveelheid energie. De grootste jaarlijkse groei vinden we echter terug bij communicatienetwerken, wat niet verwonderlijk is gezien de explosief toegenomen populariteit van mobiele communicatie.

Terugdringen van het elektriciteitsverbruik in kernnetwerken De bovenstaande resultaten pleiten er sterk voor om de inspanning met betrekking tot het reduceren van het ICT elektriciteitsverbruik te verhogen, en zeker in kernnetwerken. Kernnetwerken zijn, simpel gesteld, de grote verborgen 'snelwegen' van het Internet. Deze netwerken bestaan uit o.a. grote Internet Protocol (IP) routers die onderling gelinkt zijn via hoge-capaciteits glasvezelverbindingen. De grootte van deze netwerken kan variëren maar strekt zich vaak uit over de oppervlakte van een land. Echter, naast ecologische motieven, zijn er nog twee belangrijke redenen om het elektriciteitsverbruik in voornamelijk kernnetwerken te reduceren. Ten eerste, gezien het merendeel van deze apparatuur geconcentreerd zit in computerruimtes van netwerkoperatoren, brengt dit belangrijke technische uitdagingen met zich mee wat betreft het afvoeren van de geproduceerde warmte; wat niet zelden resulteert in hoge koelingskosten. Ten tweede, sinds ongeveer het begin van dit millennium stijgt de elektriciteitsprijs, wat een breuk is met de daaraan voorafgaande trend waarbij elektriciteit steeds goedkoper werd (gecorrigeerd voor inflatie). Beide factoren hebben een sterke impact op telecom operatoren en datacenter operatoren, gezien de elektriciteitskost een aanzienlijk deel van de operationele kosten bedraagt.

Deze thesis levert drie afzonderlijke bijdragen met betrekking tot het onderzoek naar het reduceren van het elektriciteitsverbruik in kernnetwerken. Een eerste bijdrage is een lijst met representatieve waarden voor het stroomverbruik van apparatuur in kernnetwerken; deze waarden kunnen gebruikt worden in toekomstig onderzoek. De belangrijkste drijfveer hiervoor was de observatie dat verschillende studies ook erg verschillende aannames maken naar het stroomverbruik van gelijkaardige apparatuur; dit komt een vergelijkbare evaluatie niet ten goede. De brondata is beschikbaar gemaakt in een afzonderlijk rapport en in een online database (http://powerlib.intec.ugent.be).

Een tweede bijdrage evalueert circuit geschakelde netwerken—waarbij de energie-intensieve Internet Protocol (IP) routers worden overbrugd—als een manier om het elektriciteitsverbruik te reduceren. Onze resultaten tonen dat circuit geschakelde netwerken steeds zuiniger zijn dan pakket geschakelde netwerken wanneer de gemiddelde knoop-tot-knoop traffiekvraag groter is dan de helft van de lijnsnelheid. Pakket geschakelde netwerken kunnen echter zuiniger zijn wanneer de traffiekvraag kleiner is dan de helft van de lijnsnelheid. Een belangrijke conclusie is ook dat de verhouding tussen de gemiddelde traffiekvraag en de lijnsnelheid een kritische factor is bij toekomstig, gerelateerd onderzoek.

Een derde bijdrage is een kwantitatief, samenvattend onderzoek naar de verschillende besparingstechnieken in kernnetwerken. Dit onderzoek komt tegemoet aan de nood van zowel de onderzoeksgemeenschap als netwerkoperatoren om een duidelijk overzicht te hebben over welke technieken het grootste besparingspotentieel bevatten, welke minder, en wat het totale gecombineerde besparingspotentieel is. Het besparingspotentieel van technieken die eenmalig kunnen worden toegepast bedraagt een factor 2.3x in een *Moderate Effort* (Gematigd) scenario, en een factor 31x in een *Best Effort* (Best Mogelijk) scenario. Wanneer we ook rekening houden met de historische en geprojecteerde jaarlijkse efficiëntieverbeteringen ("De Wet van Moore") dan grofweg verdubbelt dit potentieel in een tijdspanne van 10 jaar. Het grootste besparingspotentieel zit in efficiëntere koeling en stroomvoorziening, en het gebruik van slaapmodi in overgedimensioneerde apparatuur.

Terugdringen van de uitstoot van broeikasgassen in een gedistribueerd datacenter, met behulp van de zon en de wind Vanuit ecologisch oogpunt is het probleem van het toenemend elektriciteitsverbruik van ICT vooral gelinkt aan de uitstoot van broeikasgassen tijdens de productie van elektriciteit. Een alternatieve aanpak is daarom het gebruik van 'groene' elektriciteit waarvan de uitstoot van broeikasgassen veel lager is; dit is het geval wanneer elektriciteit bijvoorbeeld is opgewekt door middel van zonne-energie of windenergie. Zeker in de context van datacenters is dit een interessante piste, gezien datacenters gebouwd kunnen worden op locaties waar alternatieve energie ruim voorhanden is. Daarenboven, indien we een gedistribueerd datacenter beschouwen waar de verschillende sites op grote geografische afstand van elkaar liggen, dan kunnen data en rekentaken verhuisd worden naar die sites waar op dat momenten voldoende groene elektriciteit beschikbaar is. Dit is een zogenaamd Follow The Sun/-Follow The Wind (FTSFTW) scenario, ofwel Volg-De-Zon/Volg-De-Wind.

Deze thesis evalueert zo een FTSFTW scenario, waarbij rekening wordt gehouden met de uitstoot van broeikasgassen verbonden aan zowel het vervaardigen als het gebruik van een datacenter. De uitstoot van broeikasgassen bij het vervaardigen blijkt niet verwaarloosbaar. In een scenario waarbij extra datacenters worden gebouwd om optimaal van hernieuwbare energie gebruik te kunnen maken (gezien de geografische spreiding), kan de uitstoot lichtjes verminderd worden wanneer de geschikte condities vervuld zijn. In een gedistribueerd datacenter waarbij de nominale belasting fel onder de maximale belasting ligt, zijn er echter grote reducties in uitstoot mogelijk. Een belangrijk factor is de uitstoot van broeikasgassen van de reguliere elektriciteit; in landen en regio's waar de elektriciteit al relatief proper is zal er geen baat zijn bij het gebruik van een FTSFTW scenario in datacenters.

Summary

In the last few decades there has been a growing (scientific) consensus that the impact of humanity on our natural ecosystems is increasing beyond the bounds of sustainability. Our ecological footprint is now estimated at 1.5 planets, which means that it now takes the Earth one year and six months to regenerate what we use in a year. Clearly we are tapping into natural resources—such as water, wood, clean air, and energy—that took many hundreds, thousands or even more years to build up. While each of these resources requires our full attention to research how we can drastically reduce their usage, this book focuses on the electricity consumption of Information and Communication Technology (ICT). ICT is a broad term to describe electronic goods and technology related to computing, data processing and information transfer.

In the first part of this work we estimate the worldwide electricity consumption of ICT equipment. Subsequently we focus on reducing the electricity consumption in ICT backbone networks, and finally we evaluate the use of solar and wind power to reduce the carbon footprint of data centers.

Estimating the power consumption of ICT The starting point for this dissertation was the observation in an earlier study that the electricity consumption of ICT equipment is increasing rapidly year after year. The rate by which the combined electricity consumption of communication networks, personal computers and data centers increases was around 4% per year (doubling every 9 years) in 2008, which is faster than the 3% annual growth of the global electricity consumption. As the electricity consumption of these three categories was estimated to amount to 4% of the global electricity consumption in the same year, this would imply that by 2040 they would consume as much electricity as half the total worldwide human electricity consumption in 2008. From an environmental point of view, this growth is clearly not sustainable.

A first contribution in this dissertation is an assessment whether the yearly growth in ICT electricity consumption has changed since the 2008 study, given the increased worldwide attention for energy efficiency in all sectors, including ICT. We estimated the electricity consumption of communication networks, personal computers and data centers in 2007 and 2012. Our results indicate that the combined growth across these three categories is now 7% per year; this is lower than the 10% per year observed before 2007 for the same three categories. So indeed, ICT electricity consumption is now growing at a rate which is smaller than before 2007. An important reason for this slowdown is the shift to more efficient technologies: from desktops to laptops, from bulky CRT monitors to LCD monitors, implementing server virtualization and more efficient cooling in data centers. While this is in a sense good news, we should not be blind to the fact that it is still increasing faster than the global human electricity consumption. All three categories—communication networks, personal computers, and data centers—consume roughly an equal amount of energy. The highest growth rates are observed in telecommunication networks, which is not surprising given the explosive growth of mobile communication in the last decade.

Reducing the energy consumption in backbone networks The observations above make a strong case for intensifying the efforts to reduce the ICT energy consumption, and specifically that in backbone communication networks. Backbone networks are, simply put, the big hidden networks that act as 'highways' for the Internet. These networks consist of (among other things) large Internet Protocol (IP) routers, which are linked together through high-capacity optical fiber connections. Such networks can differ considerably in size, but they easily span the size of a country. However, apart from ecological motivations, there are two additional important drivers for reducing the electricity consumption in backbone networks specifically. First, because most backbone network equipment (such as IP routers) is densely concentrated in telecom operator buildings, which presents major technical challenges with respect to proper heat dissipation, and induces high cooling costs. Second, since the beginning of this millennium the price per unit of electrical energy has started to rise, breaking with the preceding trend where electricity was getting cheaper each year (trend numbers corrected for inflation). This has strong implications for those businesses where electricity consumption is a significant operational cost, such as telecom operators and data center owners.

This dissertation contributes to the research on reducing electricity consumption in backbone networks in three ways. First, we present a list of representative power consumption values for backbone equipment, which can be used in current and future research. This addresses the issue where different studies use(d) widely different values for similar equipment. The underlying data for these representative values have been made publicly available in a report and an online database (http://powerlib.intec.ugent.be).

Second, this dissertation evaluates optical circuit switching (i.e., optical bypassing the power-hungry IP routers) as a means to reduce the power consumption. Our results show that circuit switching is always preferable when the average node-to-node demands are higher than half the transport linerates. However, packet switching can become preferable when the traffic demands are lower than half the transport linerates. Apart from the findings on the power saving potential, a key takeaway message is that the ratio between the average demand and the transport linerate is thus a critical factor to take into account for future, related research.

Third, we perform a quantitative survey of different power saving approaches for backbone networks. This addresses the need of the research community and network operators to have a clear overview of which approaches are most promising from an energy saving perspective, which are less promising, and what the total power reduction potential is. Our results indicate that the power reduction potential of static, once-off approaches ranges from 2.3× (Moderate Effort scenario) to 31× (Best Effort scenario). Factoring in historic and projected yearly efficiency improvements ("Moore's law") roughly doubles both saving potentials on a 10 year horizon. The largest isolated power reduction potential is available in improving the power associated with cooling and power provisioning, and applying sleep modes to overdimensioned equipment.

Reducing carbon emissions in a distributed data center, using solar and wind energy From an environmental point of view, the drawback of the growing electricity consumption of ICT is mainly linked to the associated carbon emissions. An alternative approach to reduce carbon emissions is to use electricity with a low carbon footprint (such as generated through solar or wind power). In the context of data centers this creates interesting opportunities, as data centers can be located close to those sites which are optimal for renewable power generation. In addition, if we consider a distributed data center consisting of different sites at large geographical distances, computation jobs and data can be shifted to those sites where renewable energy (sun, wind) is available at that point in time. This has been referred to as a Follow The Sun/Follow The Wind (FTSFTW) scenario.

We evaluated such a FTSFTW scenario, taking into account the carbon emissions associated with the manufacturing and operation of the distributed data center. This dissertation shows that the manufacturing carbon footprint is a non-negligible factor in this scenario, but—under certain conditions—minor carbon footprint savings are possible when deploying *additional* data center sites to fully exploit the geographic availability of renewable energy. However, larger carbon footprint savings are possible when applying the FTSFTW scenario to a distributed data center where the nominal load is far below the maximum capacity. We should note that the regional carbon emission intensity of the electricity is a critical factor; countries where the electricity production is not associated with large carbon emission will not benefit from such a FTSFTW scenario.

1

Introduction

> *"The most exciting phrase to hear in science, the one that heralds new discoveries, is not 'Eureka!' but 'That's funny...' "*
>
> –Isaac Asimov (1920–1992)

This chapter places the conducted research work in context (Section 1.1), summarizes the main contributions and outlines the structure of this dissertation (Section 1.2). It also provides an overview of the publications that were authored during this research period (Section 1.3).

1.1 General introduction to 'Green ICT'

1.1.1 The relevance of Green ICT

In the last decades, the word *green* has become a fashionable term to use. It has been applied to cars, houses, food, electricity, clothing, and many other things under the sun. It is typically used as marketing term to indicate—real or perceived—that products or services are more environmentally friendly to our ecosystem compared to the regular versions of the product or service.

A similar concept that has seen widespread adoption is that of the *(ecological) footprint*. In contrast to the term 'green', the ecological footprint tries to put an actual number on the ecological impact of products or services. It does so by measuring the associated demand for natural capital (such as bioproductive land and sea area) compared to the planet's capacity. The ecological footprint[1] concept and the associated calculation method was developed around 1994 as the PhD dissertation of Mathias Wackernagel, a Swiss-born engineer [2]. The current average human ecological footprint is already at such a level that it is not sustainable, as:

> 'today humanity uses the equivalent of 1.5 planets to provide the resources we use and absorb our waste. This means it now takes the Earth one year and six months to regenerate what we use in a year.' [3]

Green ICT vs. ICT for Green This dissertation is in the field of Information and Communication Technology (ICT), which is a broad term to describe electronic goods and technology related to computing, data processing and information transfer. It is no surprise that the above two terms 'green' and 'footprint' have also found their way into this field. *Green ICT* is generally understood as the study and practice related to reducing the footprint associated with ICT. For example, the IBCN research group at Ghent University has a research domain called 'Green ICT' [http://www.ibcn.intec.ugent.be/content/green-ict]. Its objective is, not coincidentally, 'investigating several possible ways to reduce the ICT footprint [. . .] through novel network and ICT architectures'.

At this point, it might be interesting to point out that *Green ICT* is distinct from *ICT for Green*. In the latter concept, ICT is used as a *means* to reduce the footprint across other technologies and sectors such as manufacturing, transportation, heating, and power delivery. A typical example of ICT for Green is the use of videoconferencing to reduce the amount of (air) traveling. **In this dissertation we focus on Green ICT**, and consider ICT for Green out of scope. We refer the interested reader to the extensive SMARTer2020 report [4] which explores the potential of using ICT to reduce carbon emissions in other sectors.

Different footprints of ICT A necessity for any attempt at reducing the footprint of ICT is that one should have a good idea of what comprises that footprint. In the context of ICT, the strict 'ecological footprint' definition and calculation method mentioned above is not often applied. Instead, it is common to consider different varieties of 'footprints'. For example, the manufacturing and usage of ICT equipment has an associated energy footprint. Indirectly, this energy footprint has an associated carbon footprint,

---

[1] In the context of the work you are currently reading, it is amusing to note that the concept went originally by the academic name *appropriated carrying capacity*, but it was renamed to the more accessible term *ecological footprint* in reference to a student mentioning the 'smaller footprint on his desk' of a newly delivered computer [1].

as the production and consumption induces the emission of greenhouse gases; for example, the worldwide average carbon-equivalent emissions[2] associated with one unit of electrical energy are around 500 g CO_2-eq/kWh. The energy footprint and the carbon footprint are the two most often used footprint metrics. Yet other footprints could be devised to capture the impact of ICT on our ecosystem. For example, the amount of hazardous materials contained in ICT equipment, or the water consumption associated with the production of ICT equipment. The ICT-footprint initiative [http://www.ict-footprint.com], which was initiated by the European Commission, aims to find a global consensus in the ICT industry for a common definition and measurement framework within this respect. Several existing methodologies are listed on the website of the ICT footprint initiative. However, **in this dissertation we focus almost exclusively on the electricity consumption in the use phase of ICT**; however we do consider the carbon footprint in the context of data centers in Appendix A.

The electricity consumption of ICT The often-cited study by Pickavet et al. [5] estimates that the electricity consumption of ICT during operation in 2008 represented more than 8% of the worldwide electricity consumption, and about 2.6% of the worldwide primary energy consumption[3]. Pickavet et al. break down the worldwide electricity consumption across the five ICT categories shown in Fig. 1.1: data centers, network equipment, PCs, TVs and others (such as audio equipment, printers, and copiers). As can be seen, each of these categories represent a similar share in the total power consumption. Furthermore, it was estimated in the same work that the combined power consumption was growing at about 8% per year (doubling every 9 years), much faster than the total worldwide electricity consumption (3% per year). For the network equipment category the growth is even higher at 12% per year (or doubling every 6 years). An important conclusion from that study was thus that the ICT electricity consumption is growing at an unsustainable rate, and that improvements in the energy efficiency of ICT equipment are an important issue to be addressed.

Three important drivers for reducing the electricity consumption in ICT With the above observed and projected growth rate of electricity consumption of ICT, three main motivations drive the efforts to reduce its electricity

---

[2] Carbon-equivalent emissions (CO_2-eq) express the amount of CO_2 that would have the same global warming potential when measured over a given time horizon (generally 100 years), as an emitted amount of a long-lived greenhouse gas or a mixture of greenhouse gases.

[3] Worldwide about 30% of the primary energy consumption is used for electricity production. Furthermore, the conversion from primary energy to electricity typically happens at an efficiency of only 40%.

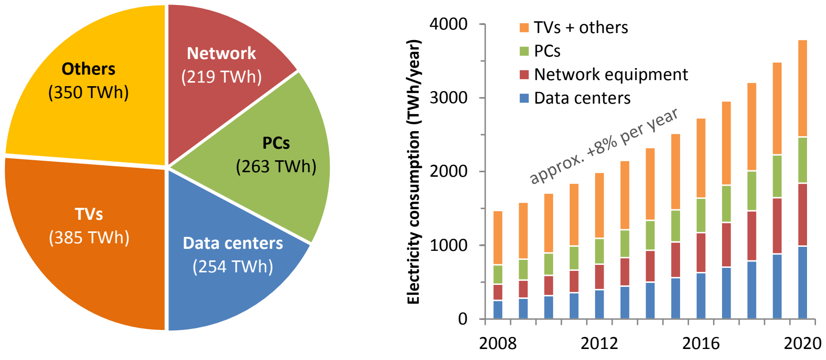

(a) **(b)**

Figure 1.1: *The worldwide electricity consumption of ICT during operation (redrawn from Pickavet 2008, [5]). (a) Situation in 2008, each category represents a similar share in the total ICT power consumption. (b) Projection from 2008 to 2020, the total ICT power consumption doubles every 9 years.*

consumption.

Economic—First, electricity prices are rising. Fig. 1.2 shows that up to the year 2005 the real price of electricity of domestic users was falling, but that from 2005 and onwards it has been rising[4] (apart from the temporary dip following the global financial crisis in 2008). Especially for those players in the ICT field where the electricity consumption is already an important part of their operational expenditures, such as data center operators and telecom operators, this has considerable impact on their business. While the rise in price has been partly caused by an increase in taxation, there has been a substantial increase in the fuel cost (gas, oil, coal, . . .) as well [6]. It is not unlikely that the increasing fuel cost is influenced by a shift from cheap conventional oil and gas to non-conventional fossil fuels (such as shale gas and extraction of oil out of tar sands) and renewables (such as solar and wind energy).

Technical—Second, the heat dissipation of ICT equipment is increasing. The ability to pack more and more transistors on an integrated circuit ("Moore's law") has led processor performance to double roughly every 18 months. However, the miniaturization has *not* been accompanied by an equal reduction in power consumption, and consequently processor power densities (in watt per surface area) have been increasing exponentially, as shown in Fig. 1.3. As a result, the power density in data processing equipment has been rising as well, as shown in Fig. 1.4. This increase is

---

[4] The trend is similar for the electricity price for industrial users [6], and is also projected to continue to rise up to 2025 [7].

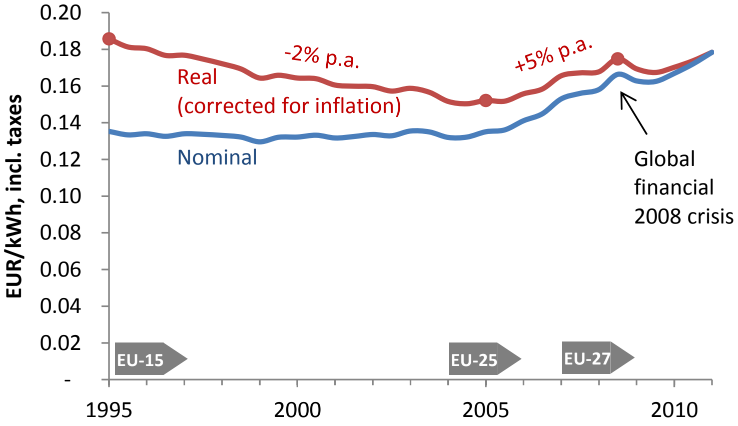

Figure 1.2: *Evolution of domestic prices of electricity (incl. taxes) in the EU from 1995 to 2011. The real price is calculated from the nominal price by correcting it with an inflation of 2% per year. The price increase since 2005 is a driver for telecom and data center operators to reduce their electricity consumption. (Source: Eurostat [8])*

creating major technical challenges to sufficiently dissipate the associate heat; especially in telecom and data centers where large amounts of such equipment are concentrated, requiring sofisticated and costly cooling solutions. According to [9], when Microsoft began charging its data center users for power consumed in response to high data center cooling costs "its users' focus changed from getting the most processing power in the smallest space to getting the most performance per watt." This is also confirmed in a 2007 report by the U.S. Environmental Protection Agency [10] which states that the situation is such that

> 'increasing power density can lead to a situation in which companies are forced to build new data centers not because they are running out of floor space but because they need power and cooling beyond what can be provided in their existing data centers. This situation has driven much of the recent interest in energy-efficiency improvements for data centers. If the power consumed (and resulting heat generated) in data centers can be reduced through energy-efficiency measures, the existing infrastructure can continue to meet cooling and power needs, and costly investments in new data centers can be deferred.'

Environmental—Third, the electricity consumption of ICT contributes to climate change. It is very probably that using fossil fuels changes the climate. The biggest contributor to climate change is the increase in the

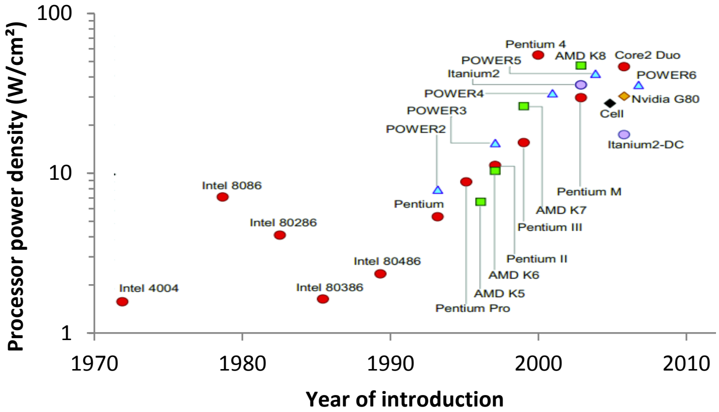

Figure 1.3: *The processor power density has been increasing exponentially. (Redrawn from Isci 2007 [11])*

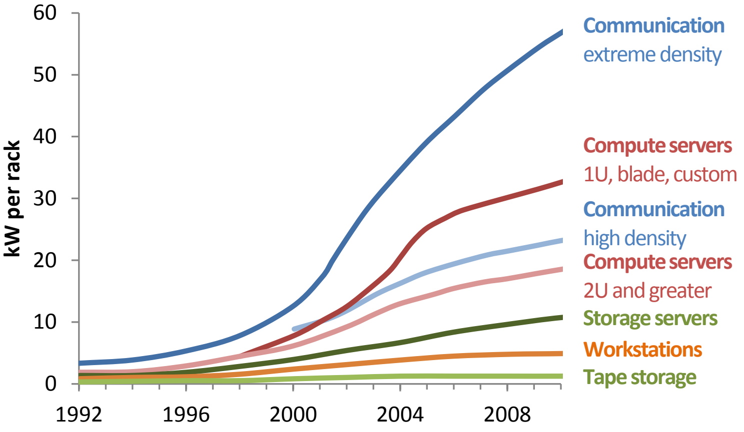

Figure 1.4: *Trends in the heat density in computer systems and telecommunications equipment. The increasing heat densities motivate data center operators and telecom operators to improve their energy efficiency, so that the cooling cost is reduced. (Redrawn from [12], based on [13] assuming rack footprint=7 sq. feet)*

greenhouse effect produced by carbon dioxide (CO_2), with most of the carbon emissions coming from burning fossil fuels.[5] As electricity is still largely generated by burning fossil fuels such as coal and natural gas (e.g., the fossil fuel contribution in the worldwide electricity generation for 2011 is 75%, for OECD countries this value is 64% [15]), electricity consumption is still an important contributor to carbon emissions[6], and thus climate change. The environmental aspect might in part be an altruistic motivation for reducing the electricity consumption associated with ICT—after all, climate change is a slow process, and the main impacts are not expected in the short term. However, initiatives, legislation, taxation and funding to reduce the use of 'dirty' (carbon intensive) electricity are getting more and more traction. The first initiative was probably the Energy Star labeling program, introduced in 1992 by the U.S. Environmental Protection Agency to promote energy efficiency in computers and related products. More recent initiatives are the guidelines on maximum power consumption for various ICT equipment, such as the EU Code of Conduct for network equipment (Fig. 1.5a), or the International Energy Agency (IEA) 1-Watt initiative to limit the power consumption of electronic equipment in standby to 1 watt (Fig. 1.5b). Furthermore, the funding for research related to Green ICT has recently been increased as well, resulting in for example EU-wide projects as ECONET (http://www.econet-project.eu/), the EU network of excellence TREND (http://www.fp7-trend.eu/), or the global GreenTouch consortium (http://www.greentouch.org/). In response, the output of research publications on this topic have increased exponentially over the last 10 years, as shown in Fig. 1.6.

1.1.2 Energy consumption in backbone networks

In the above section, we have shown that there are three important drivers for reducing the electricity consumption of ICT: economic (the cost of electricity is rising), technical (the increased heat dissipation represent technical challenges and associated costs), and environmental (the carbon footprint of the current electricity mix contributes to climate change). The combined

---

[5] While some media is eager to report on the heated debates concerning climate change, the scientific consensus is clear; an extensive discussion is outside the scope of this dissertation. For a very accessible introduction and overview, the 'Motivations' chapter in the book 'Sustainable Energy — without the hot air' by David MacKay [14] is a well-recommended, scientific *and* pleasant read, and freely available at http://www.withouthotair.com. For an in-depth analysis, see the Intergovernmental Panel on Climate Change (IPCC) assessment reports at http://www.ipcc.ch/report/ar5/.

[6] A more accurate indicator for the carbon emission contribution of electricity is the emission intensity, which is around 500 g CO_2-eq/kWh for the global average, but with large variations depending on the country or region. See Section A.4.2.

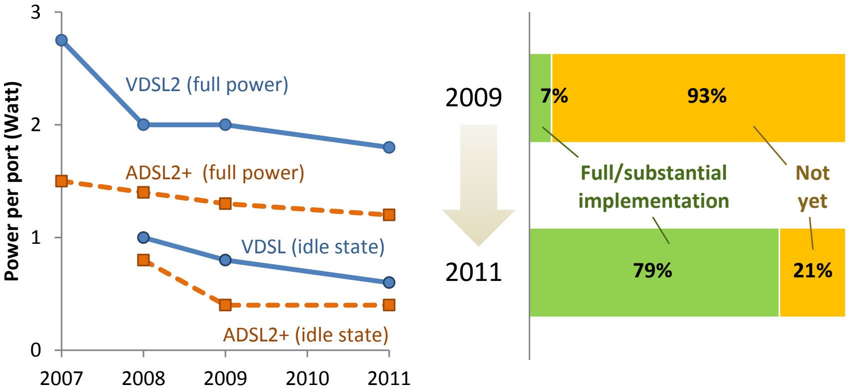

(a) Evolution of the power consumption 'Code of Conduct' guidelines for xDSL access network port values (Source: [16])

(b) Percentage of IEA member countries that implement the 1-Watt initiative for electronic equipment (consumption in standby mode limited to 1 Watt). (Source: IEA [17])

Figure 1.5: *Two examples of environmentally driven major initiatives to reduce the power consumption of ICT equipment*

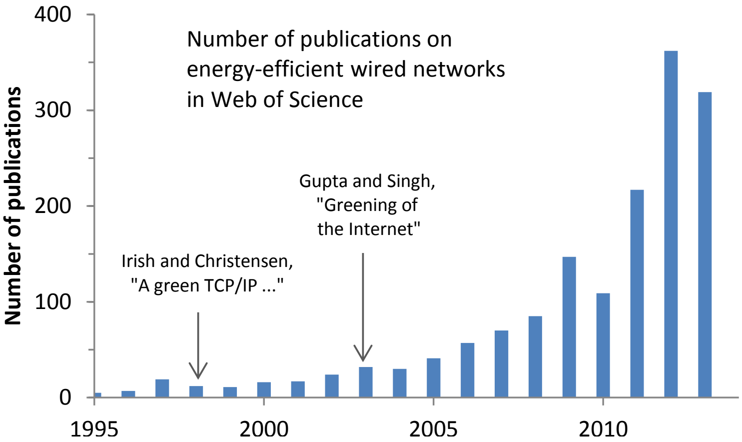

Figure 1.6: *The number of publications on energy-efficient wired networks has increased significantly in the last years. The 2003 paper by Gupta and Singh [18] is considered as a seminal work on energy-efficient networks; a number of interesting works such as by Irish and Christensen [19] precedes it however. (Source: WebOfScience with 'Topic=((green OR "energy efficient" OR "energy-efficient") AND network* NOT sensor NOT wireless NOT mobile)' and 'Research Areas=telecommunications OR computer science')*

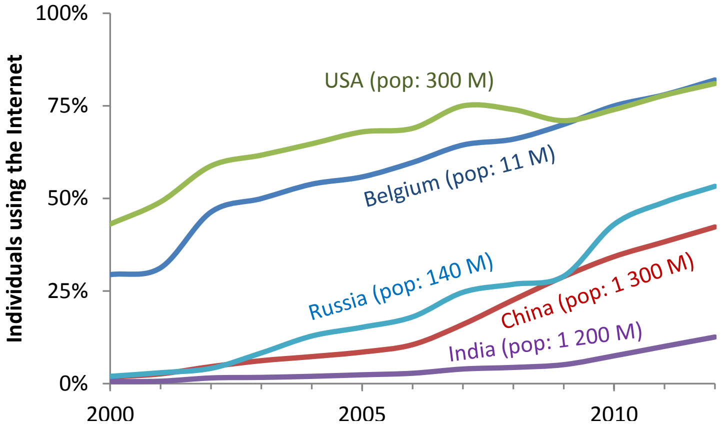

Figure 1.7: *Evolution of the national percentage of individuals using the Internet. The increasing adoption drives the growth in traffic and electricity consumption of the Internet. (Source: ITU [20])*

motivation provides a big incentive both to manufacturers, operators, governmental bodies and (to a minor extent) individual end-users to reduce the energy consumption across a large range of ICT equipment. **In this dissertation, we will mainly focus on the electricity consumption of ICT network equipment, and more specifically on backbone networks**. Before we discuss the relevance of backbone networks in the context of electricity consumption, we first give a brief primer on backbone networks.

Rise of the Internet The network-of-networks known as the 'Internet' has seen a phenomenal growth in both size and popularity from the original ARPANET around 1970, to the current global system that interconnects several billion devices worldwide. The impact of the Internet on our society can hardly be underestimated. It is probably no exaggeration to state that it has become critical from an economic point of view, and that it is significantly changing the way in which people travel, interact and relax. Fig. 1.7 shows the percentage of individuals using the Internet in a number of countries. At 80%, the adoption is almost complete in developed economies, and rising to reach similar levels in the so-called emerging economies. This drives the growth of the Internet, and the associated electricity consumption.

A breakdown of the Internet in backbone, metro and access networks As a simplification, the Internet can be split up in backbone (core) net-

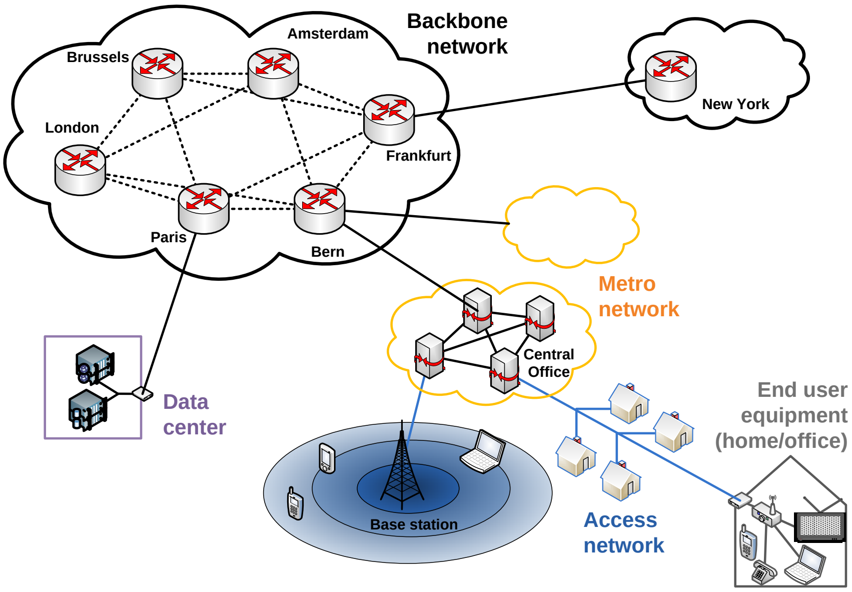

Figure 1.8: High-level and illustrative breakdown of the architecture of the Internet into backbone networks, metro networks and access networks. The focus of this dissertation is on backbone networks, and to a minor extent also on data centers.

works , metro networks and access networks, see Fig. 1.8. Internet Service Providers (ISPs) are organized in metropolitan areas, hence the name of metro networks, and they place their equipment in buildings referred to as Central Offices (COs)[7]. Almost all COs today are interconnected by optical fiber. The access network refers to the segment between a customer location and its first (serving) CO. The backbone network (also referred to as core network) interconnects the metro networks using high-capacity, low-latency connections. Because of its geographical span, and because several players operate in this field, multiple backbone networks exist, all linked together to what could be considered a single super backbone network.

Architecture of backbone networks Networks are further organized in layers, which can visually be depicted as network graphs vertically stacked on top of each other. Links in the higher layers are realized through physical connections in the lowest layer. Fig. 1.9 shows a simplified architecture of a backbone network. It consists of a number of backbone nodes that are interconnected through Wavelength Division Multiplexing (WDM) optical fiber links. The WDM fiber links carry a number of wavelengths, each hav-

---

[7] The description in this paragraph is partly based on [21].

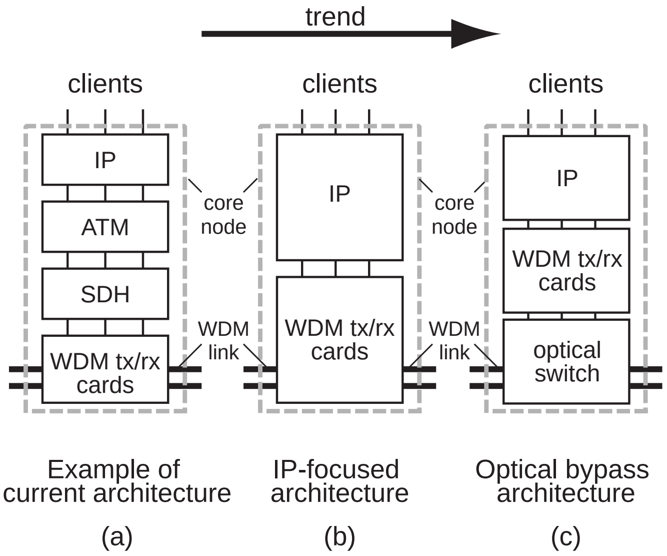

Figure 1.9: *Evolution of the architecture of backbone networks. Current node architectures consist of multiple stacked (legacy) technologies, while there is an evolution to more simplified IP-over-WDM architectures.*

ing a capacity of e.g. 10 Gbps or 40 Gbps. Optical links longer than 80 km are amplified in intermediate amplification sites (not shown). Most of the current installed equipment is typically a mix of several layers of technologies, for example an Internet Protocol (IP) packet switch on top of legacy switching technologies such as Asynchronous Transfer Mode (ATM) and Synchronous Digital Hierarchy (SDH), as shown in Fig. 1.9(a). However, there is a trend to move to more simplified architectures where IP is stacked directly on top of WDM links (Fig. 1.9(b)). As a consequence, **in this dissertation we will only focus on IP-over-WDM backbone networks**. A concept that will be centrally in this dissertation is the use of an optical switch (see Fig. 1.9(c)), where traffic not intended for the intermediate IP router remains in the optical domain and is optically switched to the next network node. As IP routers are the top energy consumers in the backbone network, this 'optical bypassing' of IP routers allows for substantial energy savings, as we will later see.

The power consumption of backbone networks The major part of the power consumption in the telecommunication operator networks is currently attributed to the wired and mobile access network. The backbone network, in contrast, is estimated to account (in 2012) for only about 8% of the total operator network consumption (which includes the wired access, mobile access, and backbone network) [22]. So at first sight there doesn't seem to be a major incentive to optimize the power consumption in backbone networks. This is probably true when considering environmental reasons as a motivation. However, the technical (and economic) drivers discussed earlier apply very much to telecom network operators, and especially to backbone networks: technical challenges and cost related to efficient cooling are a major issue in central offices and other telecom infrastructure premises. Furthermore, the energy consumption in wired access networks is proportional to the number of connected subscribers, while the consumption in the backbone network is proportional to the traffic volume [22]. With the expected increase of traffic volume, high growth rates in the backbone's energy consumption are expected, potentially even overtaking the access network's consumption [23]. The latter is illustrated in Fig. 1.10. It shows that the power per customer in the access network is more or less constant with increasing data access rates. On the other hand, the power in the backbone network increases about linearly with the data access rate.

For those reasons, it is important to react timely to the energy consumption issue of backbone networks. Hence, the motivation behind this dissertation.

1.2 Outline and research contributions

This dissertation is composed of a number of publications that were realized within the scope of this PhD. The five selected publications are so-called 'A1' publications (see Section 1.3), and provide an integral and consistent overview of the work performed. These publications have been included here in the version as they were accepted or submitted for the respective journals. In this section we explain how the different chapters are linked together, the challenges they tackle and the different research contributions following from this dissertation. The complete list of publications that resulted from this work is presented in Section 1.3.

Fig. 1.11 provides a schematic overview of all publications as a first-author (solid boxes) and a selection of publications as a co-author (dashed boxes). The arrows indicate how the findings of the various publications led to subsequent, refined publications, which culminated in the five publications (grey boxes) bundled in this dissertation.

In Section 1.1 we showed why the power consumption of telecommunication networks and ICT in general is a relevant research topic. One of the

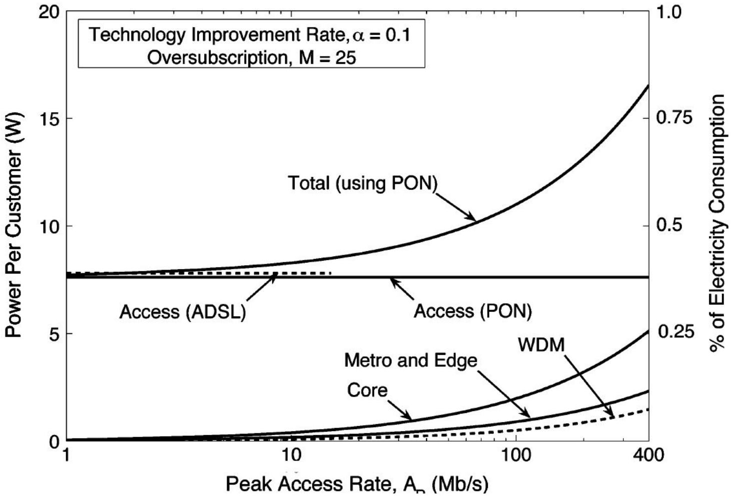

Figure 1.10: *Power consumption per customer in the access and backbone (core) network as a function of the peak access rate. According to the model, an increase in the peak access rate does not influence the power consumption in the access network, but increases the power in the backbone network linearly (note the log-scaled x-axis). (From Baliga et al. 2009, [23])*

drivers was the publication by Pickavet et al. [5] that assessed the power consumption of five major ICT categories in 2008, and showed that this was unsustainable if the associated growth trends continued in a business-as-usual scenario.

In Chapter 2 we provide an update of this earlier 'ICT footprint' study. **(Challenge 1 →)** The key research question we want to answer in this chapter is whether the worldwide power consumption of ICT in 2012 indeed increased as was projected by Pickavet et al. [5] in 2008. In addition, given the increased focus on energy efficiency and the intensified research into energy efficiency (see Fig. 1.6), can we notice any significant influence? Thereto, in Chapter 2 we consider three main ICT categories that can act as proxies for the trend of all ICT categories, and we estimate their electricity consumption and growth from 2007 to 2012. We consider (a) communication networks, (b) personal computers, and (c) data centers. **(Contribution 1 →)** As a first contribution in this dissertation, we estimate that the combined growth in electricity consumption between 2007 and 2012 is about 7% per year (i.e. doubling every 10 years); this is lower than the 10% per year observed before 2007 for the same three categories by Pickavet et al. [5]. An important reason for this decrease is a shift to more energy-efficient technologies (i.e., desktops to laptops, CRT monitors to LCD monitors, and implementation of virtualization and more efficient cooling in data centers).

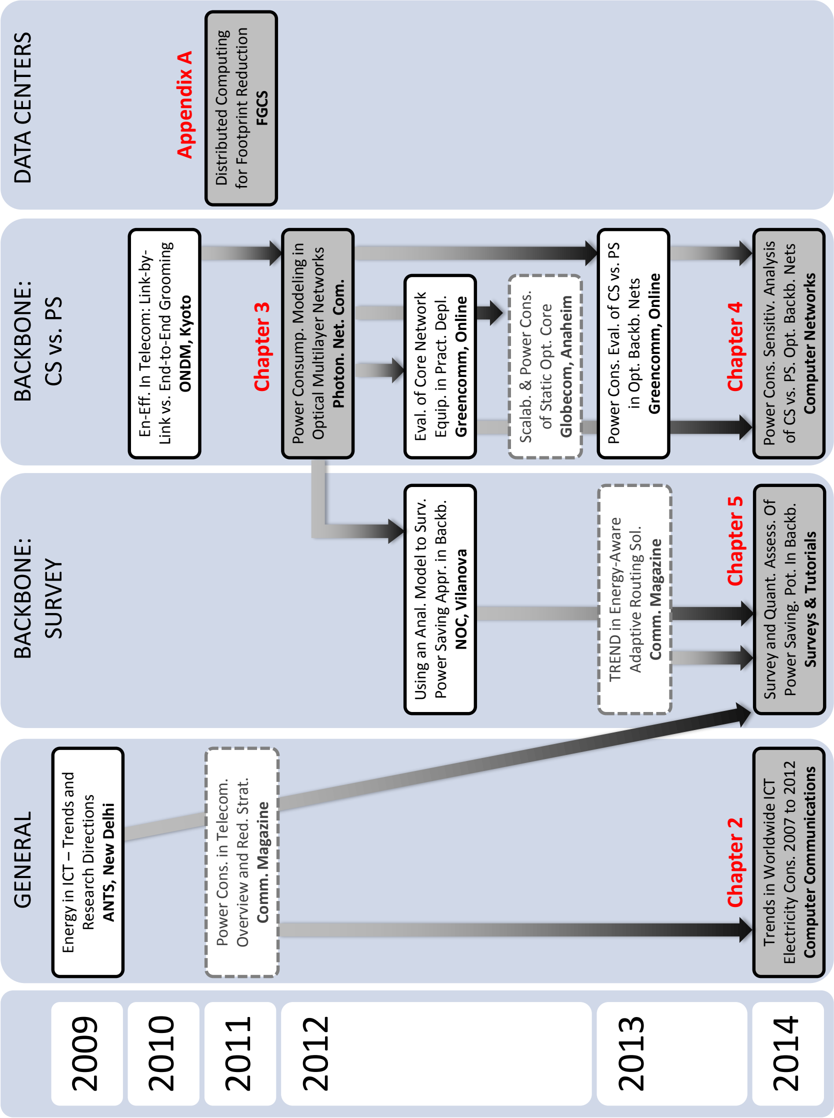

Figure 1.11: *Schematic overview of the main publications and their relation to the five main chapters in this dissertation. Solid boxes indicate first-author publications; dashed boxes indicate relevant publications with significant contribution as a co-author. The arrows indicate how the findings of the various publications led to subsequent publications.*

Together, these three ICT categories consume in 2012 about 4.6% of the worldwide electricity consumption. A rough estimate of the remaining ICT equipment (such as TVs, set-top boxes, DVD-players, and mobile phones) bumps up this figure to about 9%.

In Chapter 3 we zoom in on the electricity consumption of backbone networks. An additional observation from our above footprint study is that the electricity consumption of communication networks is growing faster than the other two categories. This implies that research into reducing the power consumption of these networks is still very relevant. **(Challenge 2 →)** With respect to existing publications focusing on backbone networks, we noticed three issues concerning the equipment power consumption values used for research in improving the energy efficiency in backbone networks. First, these values differ considerably across different studies; for example, the power consumption of an inline optical amplifier is taken 8 W in [24] and 1000 W in [25]. Second, the scope of the power consumption values is not always clear and consistent; sometimes the power consumption value covers only the basic functionality, at other times it includes system overhead (like chassis fans), and again in other cases the value might be either a typical value or on the other hand a maximum value intended for power provisioning the infrastructure. Finally, references to equipment datasheets often are overly general (e.g., the home page of the vendor's website) or become obsolete (e.g., consider the references *[25]* and *[27]* in [26]). **(Contribution 2a →)** Therefore, our (first) contribution to power consumption studies in backbone networks is to create a representative set of power consumption values for the various equipment in IP-over-WDM backbone networks. Thereto, we collected a large number of equipment data sheets and power consumption references, and made these available in a separate report [27]. From this we distill representative values for each equipment type, which are subsequently used in our further studies, and notably that in Chapter 4. **(Contribution 2b →)** In addition, the study also provides a simple analytical model to estimate the power consumption in a backbone network; this model will be used in Chapter 5 as a basis for our survey of the total power saving potential in backbone networks.

(Challenge 3 →) In Chapter 4 we are interested in the power consumption trade-off between (optical) circuit switching versus (electronic) packet switching. An observation from our work in Chapter 3, and also noted by many other works, is that IP routers are the most power hungry components in an IP-over-WDM backbone network. Bypassing these IP routers by keeping the traffic in the energy-efficient optical layer and setting up optical circuits from end-to-end thus seems an attractive solution. On the other hand, the IP packet switches allow to groom (i.e., aggregate) traffic in intermediate network nodes, thereby maximally utilizing the capacity of the optical channels and thus improving the energy efficiency when traffic demands are relatively low. There are a number of studies (including our work in Chapter 3) where circuit switching[8] is proposed as a way to reduce the power consumption in backbone networks by up to 50%. However, the reported power saving potential is not a silver bullet solution, but depends on the assumptions with respect to the network topology, transport architecture, and demands. In Chapter 4 we analyze the trade-off between circuit switching and packet switching with respect to a number of parameters, notably the network mesh degree (a metric for the connectivity in the network), and the ratio between the average demand and the transport linerate. **(Contribution 3 →)** Our results show that, in general, circuit switching is preferable. Circuit switching is always preferable when the average node-to-node demands are higher than half the transport linerates. However, packet switching can become preferable when the traffic demands are lower than half the transport linerate. We find that an increase in the network node count does not consistently increase the energy savings of circuit switching over packet switching, but is heavily influenced by the mesh degree and (to a minor extent) by the average link length. Another key take-away message is that the ratio between the average demand and the transport linerate has considerable effect on the overall efficiency, and is a critical factor to take into account for future, related research.

In Chapter 5 we use the analytical model from Chapter 3 to survey the combined power saving potential in backbone networks of various energy-efficient techniques. **(Challenge 4 →)** While over the last years a considerable number of publications have appeared proposing various techniques to reduce the power consumption in backbone networks, so far no clear assessment has been made on the total power saving potential when combining these various approaches. This assessment is not straightforward, as (a) some of these approaches focus on specific equipment only (e.g. the WDM layer), (b) certain approaches might reduce the potential of other power reduction techniques, and (c) the various publications use different baselines (topologies, traffic demands, architectures, equipment power values ...) to evaluate the power saving potential. Therefore we are interested in assessing both the combined power reduction potential of the various power saving approaches proposed in the existing body of research, as well as the relative share of each individual contribution. This way, the most promising approaches can be identified. In our work described in Chapter 5 we first describe how our analytical model can be used as a framework to map different approaches to different contribut-

---

[8] In this context, circuit switching is also referred to as 'optical bypass'

ing factors to the power consumption in a backbone network. We then survey a large number of existing research on backbone power saving techniques, and derive a Moderate Effort and a Best Effort reduction factor for a number of major techniques. **(Contribution 4 →)** Our estimates indicate that the combined power reduction potential of the once-only approaches is a factor 2.3× in the Moderate Effort scenario and 31× in a Best Effort scenario. Factoring in the historic and projected yearly efficiency improvements ("Moore's law") roughly doubles both values on a 10 year horizon. The largest isolated power reduction potential is available in improving the power associated with cooling and power provisioning, and applying sleep modes to overdimensioned equipment.

The work in Appendix A has a slightly different focus compared to the earlier chapters; i.e., we shift our focus from backbone networks to data centers. In addition, we broaden our footprint scope from electricity consumption to carbon emissions. Next to improving the energy efficiency, the use of renewable energy has been proposed as a critical measure to reduce the rising carbon emissions associated with ICT. However, the use of, and shift to, renewable energy is not without issues. The power supply by renewable energy sources such as solar and wind is typically intermittent in time. Furthermore, good sites for such renewable installations are often located in distance areas, with the power infrastructure not yet deployed or suitable for dealing with these new flow patterns. Data centers, on the other hand, can be built close to such renewable power installations, whereby data rather than power is transported out of these locations. In addition, by dynamically shifting jobs and data to sites where renewable power is currently available, such a distributed data center could operate at a minimum carbon footprint. Such a scenario has been referred to as a Follow The Sun/Follow The Wind (FTSFTW). **(Challenge 5 →)** Thus, in Appendix A we want to evaluate under which conditions such a FTSFTW scenario can indeed reduce the overall carbon footprint of data centers, and to which extent. To do so we devise a mathematical model for calculating the carbon footprint and savings of such a distributed data center, which is powered by a fixed mix of low-footprint (e.g., wind generated) and high-footprint (e.g., coal generated) energy. **(Contribution 5 →)** Our contribution is that we show that the manufacturing carbon footprint is a non-negligible factor in footprint reductions, but that—under certain conditions—minor carbon footprint savings are possible when deploying *additional* data center sites to fully exploit the geographic availability of renewable energy. However, larger footprint savings are possible when applying the FTSFTW scenario to a distributed data center where the nominal load is well below the maximum capacity.

1.3 Publications

The research results obtained during this PhD research have been published in scientific journals and presented at a series of international conferences. The following list provides an overview of the publications during my PhD research.

1.3.1 Publications in international journals (listed in the Science Citation Index [9])

1. Willem Vereecken, **Ward Van Heddeghem**, Margot Deruyck, Bart Puype, Bart Lannoo, Wout Joseph, Didier Colle, Luc Martens, and Piet Demeester, *Power Consumption in Telecommunication Networks: Overview and Reduction Strategies*, In Communications Magazine, IEEE, 49(6):62–69, June 2011. doi:10.1109/MCOM.2011.5783986

2. **Ward Van Heddeghem**, Willem Vereecken, Didier Colle, Mario Pickavet, and Piet Demeester, *Distributed Computing for Footprint Reduction by Exploiting Low-Footprint Energy Availability*, In Future Generation Computer Systems, 28(2):405–414, Feb. 2012. doi:10.1016/j.future.2011.05.004

3. **Ward Van Heddeghem**, Filip Idzikowski, Willem Vereecken, Didier Colle, Mario Pickavet, and Piet Demeester, *Power consumption modeling in optical multilayer networks*, In Photonic Network Communications, 24(2):86–102, Oct. 2012. doi:10.1007/s11107-011-0370-7

4. Sofie Lambert, **Ward Van Heddeghem**, Willem Vereecken, Bart Lannoo, Didier Colle, and Mario Pickavet, *Worldwide Electricity Consumption of Communication Networks*, In Optics Express, 20(26):B513–B524, Dec. 2012. doi:10.1364/OE.20.00B513

5. Filip Idzikowski, Edoardo Bonetto, Luca Chiaraviglio, Antonio Cianfrani, Angelo Coiro, Raúl Duque, Felipe Jiménez, Esther Le Rouzic, Francesco Musumeci, **Ward Van Heddeghem**, Jorge López Vizcaíno, and Yabin Ye, *TREND in Energy-Aware Adaptive Routing Solutions*, In Communications Magazine, IEEE, 51(11):94–104, Nov. 2013. doi:10.1109/MCOM.2013.6658659

---

[9] The publications listed are recognized as 'A1 publications', according to the following definition used by Ghent University: A1 publications are articles listed in the Science Citation Index Expanded, the Social Science Citation Index or the Arts and Humanities Citation Index of the ISI Web of Science, restricted to contributions listed as article, review, letter, note or proceedings paper.

6. **Ward Van Heddeghem**, Sofie Lambert, Bart Lannoo, Didier Colle, Mario Pickavet, and Piet Demeester, *Trends in Worldwide ICT Electricity Consumption from 2007 to 2012*, In Computer Communications, 50:64–76, Sept. 2014. doi:10.1016/j.comcom.2014.02.008
7. Yabin Ye, Felipe Jiménez Arribas, Jaafar Elmirghani, Filip Idzikowski, Jorge López Vizcaíno, Paolo Monti, Francesco Musumeci, Achille Pattavina, and **Ward Van Heddeghem**, *Energy Efficient Resilient Optical Networks: Challenges and Trade-offs*, Submitted to Communications Magazine, Feb. 2014.
8. **Ward Van Heddeghem**, Filip Idzikowski, Francesco Musumeci, Achille Pattavina, Bart Lannoo, Didier Colle, and Mario Pickavet, *A Power Consumption Sensitivity Analysis of Circuit-Switched Versus Packet-Switched Backbone Networks*, Accepted for publication in Computer Networks, Sept. 2014.
9. **Ward Van Heddeghem**, Bart Lannoo, Didier Colle, Mario Pickavet, and Piet Demeester, *A Quantitative Survey of the Power Saving Potential in IP-over-WDM Backbone Networks*, Submitted to Communications Surveys and Tutorials, IEEE, May 2014.

1.3.2 Publications in other international journals

None

1.3.3 Publications in book chapters

1. Bart Puype, **Ward Van Heddeghem**, Didier Colle, Mario Pickavet, and Piet Demeester, *"Energy-efficient Traffic Engineering"*. Chapter in the book *Cross-layer design in Optical Networks 15:199–222*, ISBN: 978-1-46145-670-4, Springer, 2013. doi:10.1007/978-1-4614-5671-1_10

1.3.4 Publications in international conferences (listed in the Science Citation Index [10])

1. **Ward Van Heddeghem**, Willem Vereecken, Mario Pickavet, and Piet Demeester, *Energy in ICT - Trends and Research Directions*, In Advanced Networks and Telecommunication Systems (ANTS), 2009 IEEE 3rd

---

[10] The publications listed are recognized as 'P1 publications', according to the following definition used by Ghent University: P1 publications are proceedings listed in the Conference Proceedings Citation Index - Science or Conference Proceedings Citation Index - Social Science and Humanities of the ISI Web of Science, restricted to contributions listed as article, review, letter, note or proceedings paper, except for publications that are classified as A1.

International Symposium on, pages 1–3, New Delhi, India, 2009. doi:10.1109/ANTS.2009.5409881

2. Willem Vereecken, **Ward Van Heddeghem**, Bart Puype, Didier Colle, Mario Pickavet, and Piet Demeester, *Optical Networks: How Much Power Do They Consume and How Can We Optimize This?*, In Optical Communication (ECOC), 2010 36th European Conference and Exhibition on, pages 1–4, Torino, Italy, 2010. doi:10.1109/ECOC.2010.5621575

3. **Ward Van Heddeghem**, Filip Idzikowski, Esther Le Rouzic, Jean Yves Mazeas, Hubert Poignant, Suzanne Salaun, Bart Lannoo, and Didier Colle, *Evaluation of Power Rating of Core Network Equipment in Practical Deployments*, In Online Conference on Green Communications (GreenCom), 2012 IEEE, pages 126–132, Sept. 2012. doi:10.1109/GreenCom.2012.6519628

4. Michael C. Parker, Richard Martin, Stuart D. Walker, **Ward Van Heddeghem**, and Bart Lannoo, *Energy-Efficient Master-Slave Edge-Router Upgrade Paths in Active Remote Nodes of Next-Generation Optical Access*, In Asia Communications and Photonics Conference, OSA Technical Digest (online) (Optical Society of America, 2012), paper ATh1D.5., Guangzhou China, Nov. 2012. doi:http://dx.doi.org/10.1364/ACP.2012.ATh1D.5

5. Slavisa Aleksic, **Ward Van Heddeghem**, and Mario Pickavet, *Scalability and Power Consumption of Static Optical Core Networks*, In Global Communications Conference (GLOBECOM), 2012 IEEE, pages 3465–3471, Anaheim, USA, Dec. 2012. doi:10.1109/GLOCOM.2012.6503651

6. Esther Le Rouzic, Edoardo Bonetto, Luca Chiaraviglio, Frederic Giroire, Filip Idzikowski, Felipe Jiménez, Christoph Lange, Julio Montalvo, Francesco Musumeci, Issam Tahiri, Alessandro Valenti, **Ward Van Heddeghem**, Yabin Ye, Andrea Bianco, and Achille Pattavina, *TREND Towards More Energy-efficient Optical Networks*, In Optical Network Design and Modeling (ONDM), 2013 17th International Conference on, pages 211–216, Brest, France, April 2013.

7. Esther Le Rouzic, Raluca-Maria Indre, Luca Chiaraviglio, Francesco Musumeci, Achille Pattavina, Jorge López Vizcaíno, Yabin Ye, **Ward Van Heddeghem**, Andrea Bianco, Edoardo Bonetto, Michela Meo, Felipe Jiménez, Filip Idzikowski, and Ruben Cuevas, *TREND Big Picture on Energy-Efficient Backbone Networks*, In Digital Communications - Green ICT (TIWDC), 2013 24th Tyrrhenian Inter-

national Workshop on, pages 1–6, Genoa, Italy, Sept. 2013. doi:10.1109/TIWDC.2013.6664209

8. **Ward Van Heddeghem**, Francesco Musumeci, Filip Idzikowski, Achille Pattavina, Bart Lannoo, Didier Colle, and Mario Pickavet, *Power Consumption Evaluation of Circuit-Switched Versus Packet-Switched Optical Backbone Networks*, In Online Conference on Green Communications (OnlineGreenComm), 2013 IEEE, Oct. 2013. doi:10.1109/OnlineGreenCom.2013.6731029

1.3.5 Publications in other international conferences

1. **Ward Van Heddeghem**, Maarten De Groote, Willem Vereecken, Didier Colle, Mario Pickavet, and Piet Demeester, *Energy-Efficiency in Telecommunications Networks: Link-by-Link versus End-to-End Grooming*, In Optical Network Design and Modeling (ONDM), 2010 14th Conference on, pages 1–6, Kyoto, Japan, 2010. doi:10.1109/ONDM.2010.5431570

2. Willem Vereecken, **Ward Van Heddeghem**, Didier Colle, Mario Pickavet, and Piet Demeester, *Overall ICT Footprint and Green Communication Technologies*, In Communications, Control and Signal Processing (ISCCSP), 2010 4th International Symposium on, pages 1–6, Limassol, Cyprus, 2010. doi:10.1109/ISCCSP.2010.5463327

3. Willem Vereecken, **Ward Van Heddeghem**, Didier Colle, Mario Pickavet and Piet Demeester, *The Environmental Footprint of Data Centers: The Influence of Server Renewal Rates on the Overall Footprint*, In Proceedings of the International Conference on Green Communications and Networks (GCN 2011), pages 823–831, Chongquing, China, 2011. doi:10.1007/978-94-007-2169-2_98

4. Kim Khoa Nguyen, Mohamed Cheriet, Mathieu Lemay, Bill St. Arnaud, Victor Reijs, Andrew Mackarel, Pau Minoves, Alin Pastrama, and **Ward Van Heddeghem**, *Renewable Energy Provisioning for ICT Services in a Future Internet*, In The Future Internet: Lecture Notes in Computer Science, pages 419–429, 2011. doi:10.1007/978-3-642-20898-0_30

5. **Ward Van Heddeghem**, Michael C. Parker, Sofie Lambert, Willem Vereecken, Bart Lannoo, Didier Colle, Mario Pickavet, and Piet Demeester, *Using an Analytical Power Model to Survey Power Saving*

Approaches in Backbone Networks, In Networks and Optical Communications (NOC), 2012 17th European Conference on, pages 1–6, Vilanova i la Geltru, Spain, June 2012. doi:10.1109/NOC.2012.6249942

6. Sofie Lambert, **Ward Van Heddeghem**, Willem Vereecken, Bart Lannoo, Didier Colle, and Mario Pickavet, *Estimating the Global Power Consumption in Communication Networks*, In Optical Communication (ECOC), European Conference and Exhibition on, OSA Technical Digest (online) (Optical Society of America, 2012), paper We.1.G.1. , Amsterdam, Netherlands, Sept. 2012. doi:10.1364/ECEOC.2012.We.1.G.1

7. Dimitris Hatzopoulos, Iordanis Koutsopoulos, George Koutitas, and **Ward Van Heddeghem**, *Dynamic Virtual Machine Allocation in Cloud Server Facility Systems with Renewable Energy Sources*, In Communications (ICC), 2013 IEEE International Conference on, pages 4217–4221, Budapest, Hungary, June 2013. doi:10.1109/ICC.2013.6655225

1.3.6 Publications in national conferences

1. **Ward Van Heddeghem**, *ICT Network Power Consumption*, In 11th UGent – FirW PhD symposium, pages 177–177, Ghent, Belgium, Dec. 2010.

2. **Ward Van Heddeghem**, *Trends in ICT Electricity Consumption*, In 14th FEA PhD symposium, Ghent, Belgium, Dec. 2013.

References

[1] W. Safire. *Footprint*. The New York Times, February 17, 2008. Last accessed: 2014-04-30. Available from: http://www.nytimes.com/2008/02/17/magazine/17wwln-safire-t.html.

[2] M. Wackernagel and W. Rees. *Our Ecological Footprint: Reducing Human Impact on the Earth*. New catalyst bioregional series. New Society Publishers, 1998.

[3] Global Footprint Network. *World Footprint*. Last accessed: 2014-04-30. Available from: http://www.footprintnetwork.org/en/index.php/GFN/page/world_footprint/.

[4] GeSI and BCG. *SMARTer2020: The Role of ICT in Driving a Sustainable Future*, 2012. Available at: http://gesi.org/SMARTer2020.

[5] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt, and P. Demeester. *Worldwide Energy Needs for ICT: The Rise of Power-aware Networking*. In 2nd International Symposium on Advanced Networks and Telecommunication Systems, 2008 (ANTS '08), pages 1–3, December 2008. doi:10.1109/ANTS.2008.4937762.

[6] *Review of European Electricity Prices*. Technical report, KEMA Consulting GmbH, November 2005. Available from: http://www.eurelectric.org/download/download.aspx?DocumentFileID=38520.

[7] P. Capros, L. Mantzos, N. Tasios, A. De Vita, and N. Kouvaritakis. *EU Energy Trends to 2030 - Update 2009*, November 2010. doi:10.2833/21664. Available from: http://ec.europa.eu/energy/observatory/trends_2030/doc/trends_to_2030_update_2009.pdf.

[8] *Evolution of Domestic Prices of Gas and Electricity*. Gas in focus, 2011. Last accessed: 2014-05-05. Available from: http://www.gasinfocus.com/en/indicator/evolution-of-domestic-prices-of-gas-and-electricity/.

[9] R. L. Mitchell. *Data Center Density Hits the Wall*. Computerworld, January 21, 2010. Last accessed: 2014-05-19. Available from: http://www.computerworld.com/s/article/9144466/Data_center_density_hits_the_wall.

[10] EPA. *Report to Congress on Server and Data Center Energy Efficiency*. U.S. Environmental Protection Agency,

http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf, 2007.

[11] C. Isci. *Workload Adaptive Power Management with Live Phase Monitoring and Prediction*. PhD thesis, Princeton University, 2007.

[12] M. Patterson and R. Jones. *ASHRAE Datacom Book Series*. White Paper #18, The Green Grid, January 2010. Available from: http://www.thegreengrid.org/~/media/WhitePapers/members/ASHRAE%20Books%20White%20Paperv11.pdf?lang=en.

[13] ASHRAE. *Datacom Equipment Power Trends and Cooling Applications - 2nd ed.*, 2005. ISBN: 978-1-936504-28-2.

[14] D. MacKay. *Sustainable Energy – Without the Hot Air*. UIT Cambridge, 2008.

[15] *IEA Sankey Diagram*. IEA, 2008. Last accessed: 2014-05-06. Available from: http://www.iea.org/Sankey/.

[16] W. Van Heddeghem, W. Vereecken, M. Pickavet, and P. Demeester. *Energy in ICT – Trends and Research Directions*. In Advanced Networks and Telecommunication Systems (ANTS), 2009 IEEE 3rd International Symposium on, pages 1–3. IEEE, 2009. doi:10.1109/ANTS.2009.5409881.

[17] S. Pasquier and A. Saussay. *Progress Implementing the IEA 25 Energy Efficiency Policy Recommendations - 2011 Evaluation*. IEA, 2012. Available from: http://www.iea.org/publications/insights/progress_implementing_25_ee_recommendations.pdf.

[18] M. Gupta and S. Singh. *Greening of the Internet*. In Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM), pages 19–26, Karlsruhe, Germany, 2003. ACM. doi:10.1145/863955.863959.

[19] L. Irish and K. J. Christensen. *A Green TCP/IP to Reduce Electricity Consumed by Computers*. In Southeastcon'98. Proceedings. IEEE, pages 302–305, Orlando, USA, April 1998. IEEE. doi:10.1109/SECON.1998.673356.

[20] *Percentage of Individuals Using the Internet*. ITU, 2012. Last accessed: 2014-05-07. Available from: http://www.itu.int/en/ITU-D/Statistics/Documents/statistics/2013/Individuals_Internet_2000-2012.xls.

[21] A. Gerber and R. Doverspike. *Traffic Types and Growth in Backbone Networks*. In Optical Fiber Communication Conference, page OTuR1. Optical Society of America, March 2011.

[22] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch. *Energy Consumption of Telecommunication Networks and Related Improvement Options*. Selected Topics in Quantum Electronics, IEEE Journal of, 17(2):285–295, April 2011. doi:10.1109/JSTQE.2010.2053522.

[23] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker. *Energy Consumption in Optical IP Networks*. Journal of Lightwave Technology, 27(13):2391–2403, July 2009. doi:10.1109/JLT.2008.2010142. Available from: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4815495.

[24] G. Shen and R. Tucker. *Energy-minimized Design for IP over WDM networks*. Journal of Optical Communications and Networking, 1(1):176–186, June 2009. doi:10.1364/JOCN.1.000176.

[25] L. Chiaraviglio, M. Mellia, and F. Neri. *Reducing Power Consumption in Backbone Networks*. In Communications, 2009. ICC'09. IEEE International Conference on, pages 1–6, Dresden, Germany, 2009. IEEE. doi:10.1109/ICC.2009.5199404.

[26] A. Coiro, M. Listanti, A. Valenti, and F. Matera. *Reducing Power Consumption in Wavelength Routed Networks by Selective Switch Off of Optical Links*. Selected Topics in Quantum Electronics, IEEE Journal of, 17(2):428–436, April 2011. doi:10.1109/JSTQE.2010.2052791.

[27] W. Van Heddeghem and F. Idzikowski. *Equipment Power Consumption in Optical Multilayer Networks – Source Data*. Technical Report IBCN-12-001-01, IBCN, University of Ghent, January 2012. Available at: http://powerlib.intec.ugent.be.

2

Trends in Worldwide ICT Electricity Consumption from 2007 to 2012

In this chapter we are interested in the growth rates of the power consumption associated with telecommunication networks, personal computers and data centers. We find that the power consumption growth rate in telecommunication networks is, at 10% per year, still significantly higher than for the other two categories. This is a strong motivation for the work in the subsequent chapters on backbone networks. We would like to explicitly point out that the work in this chapter contains significant contribution from Sofie Lambert, specifically on the power consumption analysis of communication networks and personal computers.

⋆ ⋆ ⋆

Ward Van Heddeghem, Sofie Lambert, Bart Lannoo, Didier Colle, Mario Pickavet, and Piet Demeester

Published in Computer Communications, February 2014

Abstract

Information and Communication Technology (ICT) devices and services are becoming more and more widespread in all aspects of human life. Following an increased worldwide focus on the environmental impacts of

energy consumption in general, there is also a growing attention to the electricity consumption associated with ICT equipment.

In this paper we assess how ICT electricity consumption in the use phase has evolved from 2007 to 2012 based on three main ICT categories: communication networks, personal computers, and data centers. We provide a detailed description of how we calculate the electricity use and evolution in these three categories.

Our estimates show that the yearly growth of all three individual ICT categories (10%, 5%, and 4% respectively) is higher than the growth of worldwide electricity consumption in the same time frame (3%). The relative share of this subset of ICT products and services in the total worldwide electricity consumption has increased from about 3.9% in 2007 to 4.6% in 2012. We find that the absolute electricity consumption of each of the three categories is still roughly equal. This highlights the need for energy-efficiency research across all these domains, rather than focusing on a single one.

2.1 Introduction

ICT is everywhere Information and Communication Technology (ICT) devices and services have profoundly changed the way in which humans work, travel, play and interact in the last decades. An increasing number of people earn their living working in front of a computer, and many industrial and agricultural processes have in some way become controlled or monitored by intelligent electronic devices. Many cars are now equipped with a Global Positioning System (GPS) device for easier navigation on unfamiliar roads and time-of-arrival estimation, even taking into account traffic jams and road works. Entertainment has a rising digital footprint in the form of video games, online games, and in-house intelligent workout devices. The steep rise of online social services such as Facebook (over one billion users near the end of 2012 [1]) and the continued proliferation of mobile phones show that inter-human communication and interaction are increasingly taking place via digital platforms.

There is no single metric for the ICT footprint The increase of ICT equipment has an associated growing impact on our environment. This impact comes in many forms, and is often expressed as a ‘footprint’. For example, the manufacturing and usage of ICT equipment have an associated energy footprint and CO_2 emission footprint. Pollution associated with mining for rare earth metals, and waste through improper disposal of broken or end-of-life equipment can also be considered as part of the environmental footprint. As such, depending on which aspects are taken into account, there are several methodologies to measure and determine the footprint of organizations, services and goods with respect to ICT. The ICT-footprint initiative [http://www.ict-footprint.com], which was initiated by the European Commission, aims to find a global consensus in the ICT industry for a common definition and measurement framework within this respect. Several existing methodologies are listed on the website of the ICT footprint initiative.

Footprint scope of this work In this work we only consider the *use phase electricity consumption* of a number of important ICT equipment categories. For ICT equipment, the use phase has been shown to make up a large fraction of total carbon emitted during manufacturing, usage and end-of-life activities (see e.g., in [2] and [3]), and for personal computers the survey in [4] concludes that the use phase is the dominant life cycle phase for primary energy demand. We only focus on the electricity consumption, as the use phase carbon emissions can be directly calculated from the emission intensity of the electricity (i.e., the amount of CO_2 emitted per produced kWh). We do not distinguish between the electricity source being either on-grid (i.e., from a utility provider) or off-grid (such as a remote mobile base station powered by a diesel generator or solar panels). Where applicable, we do include the electricity consumption associated with cooling and power provisioning of ICT equipment in operation. While it could be argued that this broadens the scope somewhat beyond ICT equipment, we feel this is appropriate as this overhead electricity use is directly and strongly tied to that of the ICT equipment itself. To conclude, the use phase electricity consumption is a relevant factor in the overall ICT footprint, and therefore merits a dedicated, detailed study.

The purpose of estimating ICT electricity use The relevance of estimating the worldwide ICT electricity use is twofold. A first purpose is to assess whether ICT is a significant contributor to the worldwide electricity consumption, or by extension, to the worldwide carbon emissions. Such an assessment is not limited to the current situation, but given projected growth trends, also provides insight in the evolution of the ICT electricity share in the near future. A second purpose is to assess where efforts should be concentrated in order to reduce the worldwide ICT electricity consumption. Energy-efficiency efforts can only have a meaningful impact if they are focused on those areas or categories that contribute most—or are expected to do so in the near future—to the total ICT electricity consumption.

Earlier work on worldwide ICT footprint estimation There have been a number of earlier studies that estimate the worldwide ICT electricity use. In 2008, the SMART2020 report [3] explored the potential of ICT to reduce global carbon emissions, and while doing so provided an estimation and projection of the ICT footprint itself. In the same year, some of the co-authors of this current paper also published a study [5] to estimate the worldwide electricity use and embodied energy associated with ICT equipment and services. Finally, a study by Malmodin et al. [2] that appeared in 2010, provided an estimate of the 2007 worldwide greenhouse gas emissions and operational electricity use in ICT. Incidentally, while we were finalizing our current study, the SMARTer2020 report [6] appeared at the end of the year 2012. It provides an updated version of the earlier report based on more recent data and findings. We have intentionally refrained from using data provided in [6] for our current work, in order to have an independent assessment.

Goal and contributions of this paper The main goal of our current work is to provide an update of our earlier estimates published in 2008. Even more importantly, we want to explore the trends over the last five years, i.e. from 2007 to 2012, and see if there are significant differences in growth rates compared to earlier years. We estimate the worldwide electricity consumption of communication networks (Section 2.2), personal computers (Section 2.3) and data centers (Section 2.4). An overview is given in Fig. 2.1. We consider the use phase only; the electricity used to manufacture and dispose of equipment is not included. While our initial objective was to cover again the same five categories as we did in our previous work, we do not provide a detailed estimation of the electricity consumption of the TVs category and Others category, as we did not have sufficiently reliable data available for doing so. This is in itself not a major issue, as we can assess general trends for the available three categories, as described in Section 2.5. Finally, to assess the validity of our results, we perform an extensive comparison of our findings with the aforementioned earlier works (Section 2.6).

2.2 Communication networks

In preparation of this paper, we first performed a detailed estimation of worldwide electricity use of communication networks in the time frame 2007 to 2012, which was published in [7]. In this section we only provide a summary of these results; more detailed numbers and an in-depth explanation on the methodology can be found in the cited work. An important revision is the updated office networks estimation as we discovered that

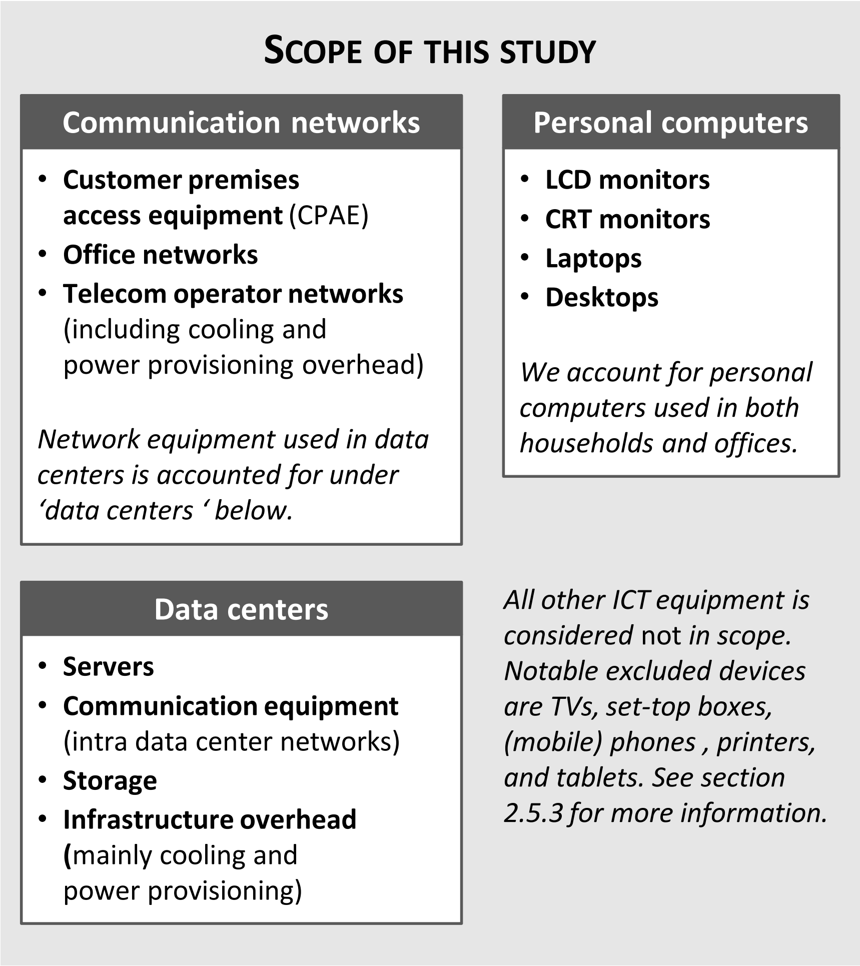

SCOPE OF THIS STUDY

Communication networks

- **Customer premises access equipment** (CPAE)
- **Office networks**
- **Telecom operator networks** (including cooling and power provisioning overhead)

Network equipment used in data centers is accounted for under 'data centers ' below.

Personal computers

- **LCD monitors**
- **CRT monitors**
- **Laptops**
- **Desktops**

We account for personal computers used in both households and offices.

Data centers

- **Servers**
- **Communication equipment** (intra data center networks)
- **Storage**
- **Infrastructure overhead (**mainly cooling and power provisioning)

All other ICT equipment is considered not *in scope. Notable excluded devices are TVs, set-top boxes, (mobile) phones , printers, and tablets. See section 2.5.3 for more information.*

Figure 2.1: *The scope of ICT equipment that we consider in this study.*

the potential double accounting for data center network equipment was larger than assumed in [7]. As a result the office network electricity use is now down to about half of our earlier estimate. We also renamed customer premises equipment to customer premises *access* equipment to make the scope less ambiguous.

We consider three components of communication networks: (a) telecom operator networks, (b) office networks and (c) customer premises access equipment. The electricity use of telecom operator networks excludes the electricity consumption in their offices and their data centers, as these are dealt with separately in Section 2.2.2 and Section 2.4; the electricity consumption of retail associated with telecom operators is considered out of scope. Electricity use in office networks includes routing, switching, WLAN, and network security equipment in offices. Finally, customer premises access equipment (CPAE) covers residential access equipment, which consists mainly of modems and routing equipment with or without Wi-Fi functionalities.

In the text below, we first describe the methodology used to assess the electricity consumption of the three subcategories, before presenting the results and a discussion on the reliability of our estimates.

2.2.1 Telecom operator networks

We estimate the worldwide electricity use for operator networks based on the electricity consumption of a selection of telecom providers. We extrapolate these operator specific values using worldwide numbers of mobile, fixed broadband and fixed telephone subscriptions.

Our approach differs from earlier approaches to estimate the worldwide electricity consumption, which are typically based on first determining the average electricity consumption per service per subscriber, through one of the following two general techniques. On the one hand, a *bottom-up* approach can be used, as was done by the authors in [8, 9]. They summed the power consumption of individual network components (such as routers, optical amplification systems and mobile base stations) to estimate the electricity consumption, per user or per unit of traffic, for a given service. These per-service consumption values were then multiplied with worldwide total traffic per service to get the total worldwide electricity consumption per service. On the other hand, in [2] a *top-down* approach was used: based on the aggregated power consumption data from a number of telecom operators, the authors determined the average electricity consumption per mobile subscriber and fixed subscriber. Multiplying these values with worldwide subscription numbers and summing the results provided them with an estimate for the worldwide electricity consumption in telecom operator networks. However, accurately determining a worldwide average electricity consumption per service per subscriber is not easy, because equipment is often shared across different services. For example, fixed and mobile services can use a single backbone network.

In order to circumvent the issue of assigning the power consumption of an operator to specific services, we use a subscription-based *representative sample* of operators. In our representative sample, the number of mobile, fixed broadband and fixed telephone subscriptions have the same relative ratios as the worldwide subscription numbers. The power consumption for this sample is then scaled up to a worldwide value by multiplying with a single scaling factor. This scaling factor is the ratio of worldwide subscriptions over the subscriptions covered in the sample, which, following the definition of the representative sample, is the same for mobile, fixed broadband and fixed telephone services. Using this approach, we don't have to determine the average power per user.

The use of a representative sample (as opposed to just taking a random combination of telecom operators) is required because we believe the power consumption per user for each of these services can differ significantly, and operators might have an unbalanced number of subscriptions for a particular service. For example, while China Mobile is by far the largest operator in our sample (in number of subscriptions), its focus is almost exclusively on mobile subscriptions. Not taking this unbalance into account would lead to worldwide electricity consumption values which are skewed by the power consumption per mobile subscriber, which can be very different from the (also unknown) power per fixed broadband or fixed telephony subscriber. Furthermore, incumbent operators often lease parts of their networks to other operators. This means that the number of customers connected to a network is not necessarily the same as the number of subscribers reported by the operator, making it difficult to determine the average electricity consumption per connected user. To cancel out the effect of leased lines as much as possible, we aggregate the subscriptions and electricity consumption of different operators.

One drawback inherent to our approach is that we cannot determine the relative contributions of different services (mobile, fixed broadband, and fixed telephone) or network sections (such as access, metro and core) to the total network electricity consumption, since we aggregate the electricity consumption for all services from the sample.

2.2.2 Office networks

The scope of this section is the electricity used by network equipment in offices, excluding network equipment in data centers. This includes network equipment in network operator offices but excludes equipment in the telecom network they operate (this was already accounted for in Section 2.2.1). We do not consider custom enterprise transport networks, such as those between Google or Amazon data centers; their power consumption will very likely be negligible as optical transport networks consume very little compared to other network equipment such as modems, IP routers or base stations.

We base our estimate on a study by Lanzisera et al. [10], which estimates the USA and worldwide electricity consumption of data network equipment in both residential buildings and offices. Their study focuses on IP-based network equipment only, and does not include the electricity used by power or cooling infrastructure. Their annual electricity consumption estimate is based on an average power consumption per device, and uses values for 2007 and 2008 with forecasts from 2009 through 2012, which we have adopted. We consider only the equipment relevant in office use, and we add an estimated overhead for cooling for each of the equipment categories, i.e. switches, routers, WLAN equipment and security equipment.

We discovered that in the earlier results published in [7], we underestimated the potential double accounting for data center network equipment[1]. Therefore, we re-evaluated our estimate. In [7] we only considered the following five categories from the study by Lanzisera: 10/100 Mb/s switches, 10/100/1000 Mb/s switches, small & medium routers, enterprise WLAN devices, and small & medium security appliances. In this study we include three extra equipment categories in our calculation: modular & 10G switches, large routers, and large security appliances. These three categories broaden the scope to completely cover all data center network equipment as well. To exclude telecom operator network equipment (which was already covered in Section 2.2.1), we consider only half of the electricity use reported in [10] for the modular & 10G switches and large routers. We do not have data to more accurately assess the share of operator network consumption in both categories; assuming the share to be half minimizes the potential deviation. As the sum of both categories contributes less than 10% of the total electricity use, the potential error will be small. We also subtract the data center network equipment electricity consumption (which

---

[1] In [7] we incorrectly reported the network equipment for data center volume servers to potentially account for 1.48 TWh (or only 5% of the total uncooled office network electricity consumption). However, the value given was in GW, which corresponds to almost 13 TWh instead of 1.48 TWh.

Table 2.1: *Office networks: cooling overhead factors and worldwide electricity use per type of equipment (electricity use estimates are adaptations of the values in [10]). The share of communication networks in data centers (determined in Section 2.4) is subtracted to avoid double accounting.*

	Cooling overhead 2007/2012	Electricity use, 2007 (TWh)	Electricity use, 2012 (TWh)	Note
switching - 10/100	1.38	12.7	10.7	
switching - 10/100/1000	1.38	5.4	17.5	
switching - mod./10G	1.95/1.83	3.9	4.3	a, b, c
routers - small & med.	1.75	3.5	4.2	
routers - large	1.95/1.83	1.0	0.4	a, b, c
enterprise WLAN	1.00	1.0	2.3	
security - small & med.	1.75	5.3	7.7	
security - large	1.95/1.83	2.9	4.0	a, b
Data center networks	1.95/1.83	-23.4	-28.9	b
Total		12.2	22.2	

[a] Equipment type not accounted for in [7]

[b] Power Usage Effectiveness (PUE) value for 2007 / 2012, from Table 2.3

[c] Half of the value specified in [10], to avoid counting the telecom operator share

can be derived from Table 2.3 in Section 2.4) to avoid overlap between the categories.

The details and results are given in Table 2.1. The worldwide office network equipment is estimated to consume 22 TWh in 2012 (instead of 42 TWh as reported earlier in [7]).

2.2.3 Customer Premises Access Equipment (CPAE)

In this section, we consider the electricity consumption of residential network access equipment. In order to access the network, every Internet subscriber requires a modem. Most users also have a Wi-Fi router installed, often with integrated wired switching and routing capabilities. The modem and Wi-Fi router may also come in a single box. We estimate the worldwide power consumption by multiplying average power consumption values of these residential devices per access technology category with the number of subscriptions per category. We consider the following access technologies: cable, Digital Subscriber Line (DSL), Fiber To The Home (FTTH), narrowband, and other broadband (such as satellite). Our scope does not include residential stand-alone wired switches, but they have been estimated in [10] to be only a small contributor. Similarly, we do not include power line communication devices; we estimate[2] that their consumption is around 2 TWh/y, but an in-depth study would be needed for a more accurate evaluation.

The number of worldwide users for each category is based on various sources: the worldwide average number of broadband subscriptions per 100 inhabitants [14], worldwide population data [15] and access technology distribution [16–18]. Values for 2012 are extrapolations based on data from previous years. The power consumption per user values for cable, DSL and FTTH are based on [10, 19]. For the relatively small number of users accessing the Internet through other broadband technologies we assumed a power consumption comparable to that of the more common broadband technologies. The per user power consumption for narrowband users is based on the power consumption of a dial-up modem [20].

2.2.4 Results

The total worldwide electricity consumption in communication networks grew from 200 TWh per year in 2007 to 330 TWh per year in 2012, corresponding to an annual growth rate of 10.4% (see Fig. 2.2).

Telecom operator networks power consumption makes up 77% of this value, customer premises access equipment about 15% and office networks only around 7%. The annual growth rate of office networks is highest with 12.8%, whereas the other two categories grow at a slightly lower rate of 10.2% (telecom operator networks) and 10.8% (customer premises access equipment).

It is interesting to note that a number of studies that use a bottom-up approach to estimate the communication networks electricity use, attribute a much larger relative share to customer premises access equipment than we do, e.g., [9, 21, 22]. For [9, 21], we believe this to be because of the constraints inherent to a bottom-up approach, which might not easily account for such things as legacy equipment, underutilized equipment, or unknown overhead in general. Furthermore, it is important to be aware of the considered scope of customer premises access equipment when comparing results; e.g., the ‘home networks’ category in [22] includes not only access equipment but also DECT telephones, set-top boxes and laptops and computers.

---

[2] HomePlug is the dominant standard for power line communication devices. The HomePlug Powerline Alliance reported over 60 million installed devices in 2010 [11]. With an estimated electricity consumption of 4 W per device (based on a number of data sheets, e.g. [12, 13])), this results in 2.1 TWh/y.

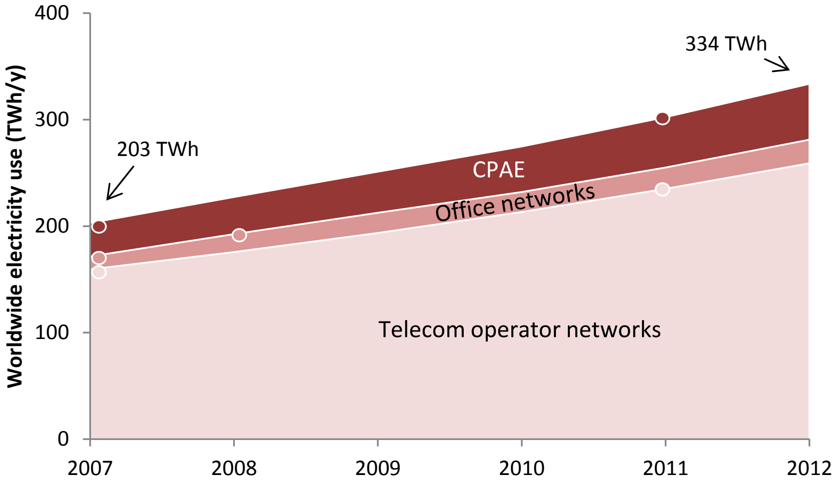

Figure 2.2: *Worldwide use phase electricity consumption of communication networks. The annual growth is 10.4% in the 2007 to 2012 time frame. Telecom operator networks dominate the result. The circular markers* ○ *indicate years for which the subcategory results are (mainly) based on data for that specific year; non-marked data points are (mainly) interpolations or extrapolations.*

2.2.5 Reliability

For the telecom operator networks calculations, we based ourselves on aggregated operator power consumption values (rather than averaging and upscaling power per subscriber values) to minimize the effect of leased and rented lines. This effect may have an influence nonetheless. Secondly, our results are—for the most part—based on publicly available electricity consumption values. As companies that publish these values are typically those that have already made efforts to improve their energy-efficiency, this may lead to overly optimistic results.

For our office networks estimations, there is some uncertainty in (a) the cooling overhead factors, as they are based on limited discussions with industrial experts, and (b) whether we accounted sufficiently for the overlap in scope with telecom operator networks. Smaller switching equipment (i.e., 10/100 and 10/100/1000 switches) will also be present in telecom operator networks, but it is unclear which fraction it represents in both switching categories in Table 2.1, and we did not account for it. This makes our estimation for office networks electricity consumption an overestimation. In any case, as office networks contribute less than 7% to the overall communication networks result, the influence of potential scope overlap on the overall result will be very small. To get an indication of the reliability of our result, it is instructive to estimate the electricity use of office network equipment per office computer, similar to what was done by Kawamoto et. al. [23]. With 548 million office computers in 2012, we get 4.6 W/unit. While this seems to be in line with the 3.8 W/unit reported by Kawamoto (in 2002), he does *not* account for cooling and does not include WLAN and security equipment in his calculation (which would result in a higher value). Malmodin et al. [2] report a value of 8 W/unit. But, as detailed in Appendix S8 of the supporting information for [2] this value includes 'faxes and other business systems' which accounts for over half of the total office network equipment consumption (or about 4 W/unit). As office end-user equipment is not in our scope, this would explain our lower overhead per office personal computer. So part of the large variation in watts per office personal computer seems to be explained by a different scope of office network equipment.

The reliability of our customer premises access equipment results depends strongly on the accuracy of our power per user estimates, which are based on averages for the USA. We were unable to the determine the evolution of average power consumption per device from 2007 to 2012. However, we do take into account shifts between different technology categories—the decrease in narrowband and increase in FTTH being the most notable—which leads us to believe the general trend in our results provides a good estimate of the evolution in power consumption of customer premises access equipment. Finally, it might be interesting to note that Lanzisera's [10] Customer Premises Access Equipment (CPAE) estimate is based on OECD countries (which currently do not include emerging economies such as Brazil, Russia, India, and China). As a result, our customer premises access equipment estimates (which do include all countries) are substantially higher.

2.3 Personal computers

The Personal Computers (PCs) category covers the electricity consumption of desktops, laptops and (external) monitors connected to computers. We exclude terminals connected to the mainframe and devices such as smartphones or tablets that have only some, but not all, of the functions of a PC (e.g. they may lack a full-sized keyboard, a large screen, ...) [24].

We calculate the worldwide energy consumption by multiplying average energy consumption values per device by numbers of devices. We distinguish between household and office desktops and laptops, and Cathode Ray Tube (CRT) and Liquid Crystal Display (LCD) monitors, as listed

Table 2.2: *Personal computers and computer monitors: average energy consumption per device (taking into account active and inactive times) and worldwide electricity use per type of equipment.*

	Energy/device (kWh/y)		Worldwide energy use (TWh)	
	2007	2012	2007	2012
Office desktops	149	137	51.4	46.2
Household desktops	231	213	91.2	105.9
Office laptops	46	39	4.1	8.3
Household laptops	70	59	17.7	45.2
Total computers			164.4	205.6
CRT monitors	175	175	46.6	31.9
LCD monitors	70	70	27.9	69.6
Total monitors			74.5	101.5
Total			238.9	307.1

in Table 2.2.

2.3.1 Number of personal computers

We estimate the worldwide number of PCs based on the average number of PCs per 100 inhabitants for each country [25] and population data for these countries [15] (we used the medium variant for population prospects). There are some gaps in the United Nations (UN) data for the number of PCs per 100 inhabitants. For some countries the data is missing for one or two years. We fill in these blanks by making a linear interpolation of the previous and the next year for which data is available. For other countries there is little or no data available, so we can't interpolate data from other years. We assume the number of PCs per 100 inhabitants in these countries equals the average value for the region they belong to. Based on these assumptions, we estimate the total number of PCs per region and worldwide for 2000-2006. From 2007 onwards, there is not enough data available in the UN database to make a reliable estimate.

However, annual PC sales numbers are available for 1991-2010 [26]. If we know the lifetime distribution of PCs, we can use these sales data to determine the number of PCs in use in 2007-2010. We model the lifetime distribution of personal computers as a curve that is initially flat, followed by an exponential decay. This curve is characterized by two parameters: the threshold and the decay constant. Based on the number of personal computers in use in 2000-2006 and the sales data for 1991-2006, we estimate the threshold of the lifetime curve is at 2.5 years, after which 26% of the PCs still in use are discarded each year. This corresponds to an average lifetime of 5.9 years. Combining this lifetime model with historical sales data (and an exponential extrapolation of this sales data to predict sales in 2011-2012) provides us with an estimate for the number of PCs in use in 2007-2012.

Based on these calculations we estimate over 1 billion personal computers were in use in 2007. We estimate this number has increased to just over 1.8 billion by the end of 2012.

2.3.2 Number of laptops and desktops, and household and office computers

Laptops typically consume much less energy than desktops. We therefore need an estimate of the number of laptops and desktops that are in use. This can be derived indirectly from the annual sales data for laptops and desktops [26, 27] and the lifetime model of personal computers we determined in the previous section. The share of laptops has known a strong increase in the past five years, from about 32% of the installed base of personal computers in 2007 to 54% in 2012.

A distinction is made between computers that are used in an office environment and computers that are used in households, since the usage patterns in these environments differ. In [19]—a study on the electricity consumption of consumer electronics in households—the number of desktops and laptops in USA households are given. Combining these numbers with the total installed base of laptops and desktops in the USA (obtained in the calculations in the previous paragraph) allows us to estimate the distribution of computers per type (laptop/desktop) and environment (household/office). We assume the worldwide distribution is similar to that in the USA.

2.3.3 External monitors

The screens integrated in laptops were already taken into account in the previous section, but we still need to consider external displays attached to most desktops and some laptops. Unfortunately we could not find any worldwide estimates for the number of computer monitors that are currently in use. In [19], survey results for the year 2010 indicate that in USA households there are on average 0.96 external displays connected to a desktop[3], and there are on average 0.26 external displays connected per laptop. We assume these fractions apply to all laptops and desktops worldwide

---

[3] It might be surprising that there is on average slightly less than one monitor per desktop computer. As the study notes: 'this is partly due to the prevalence of all-in-one PCs, i.e. those with integrated displays'.

to obtain the worldwide number of external computer monitors in use in 2010. We can't simply apply these fractions for other years as well, since the number of computer monitors per device has increased over the years. To estimate the growth rate for the number of monitors, we also use data from the USA study, where the number of computer monitors in households in 2005, 2006 and 2010 are given. Based on these numbers we expect the number of monitors to increase by 12.06% annually. We apply this growth rate to the 2010 value we obtained above to estimate the worldwide number of monitors for 2006-2012.

We make a distinction between CRT and LCD monitors, since the latter are more energy efficient. We did not find historic trends for the percentage of CRT displays in use in all regions, but we are able to derive the penetration curve of CRTs in the USA installed base from values for 2006-2010 in [19] and the fact that the first LCD monitors were commercially available around 1999 [28]. We then use the difference in transition time from CRT to LCD TVs (in sales data) as an indication for how many years we should offset the USA curve in time for other regions. For example, Indian LCD TV shipments surpassed those of CRT TVs in 2012, while the USA and Europe saw their LCD TV shipments exceed those of CRTs in 2007. This means that we shifted the curve for the percentage of CRT monitors in India 5 years into the future. Combining these curves with the installed base of computers per region provides us with a weighted average for the percentage of CRT and LCD monitors in use worldwide.

2.3.4 Power consumption per device

To the best of our knowledge, there are no *worldwide* values available for the average power consumption of desktops and laptops. One of the main challenges when determining the average power consumption of these devices is that even though the numbers for power consumption in active, sleep and off mode are known, we have no recent information on how many hours computers are left on and in sleep mode during the day. Although there are no worldwide averages available, we did find average values for the USA [19], so we derived our estimates from these numbers as follows. The work in [19] provides average per-device power consumption values for desktops and laptops in USA households for 2010, as well as a comparison to values for 2006 from a previous study. Based on these numbers we estimate the evolution in power consumption per household desktop and laptop for 2007-2012. Additionally, [19] references studies on the power consumption of office desktops and laptops, giving a value for 2005 and 2009 respectively. We assume the ratio of office to household power con-

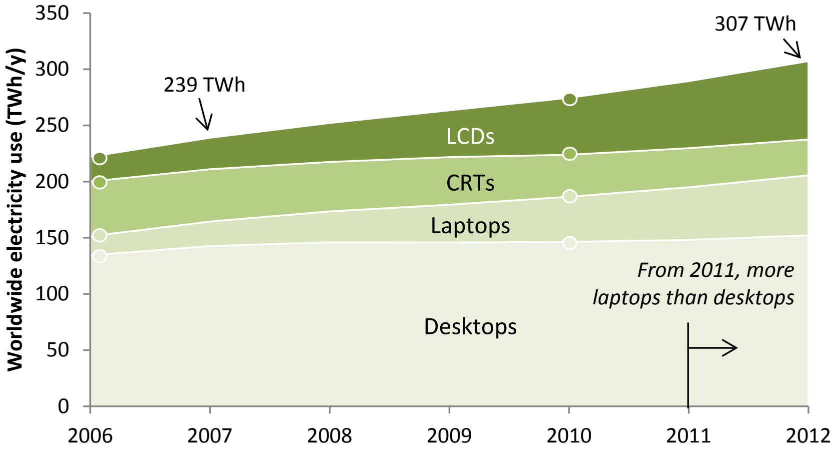

Figure 2.3: *Worldwide use phase electricity consumption of personal computers. Desktops still dominate the result. The shift to more energy-efficient technologies has tempered overall electricity use: while the number of desktops+laptops has grown at 10.9% per year from 2007 to 2012, the electricity consumption has only grown at 5.1% per year. The circular markers ○ indicate years for which the subcategory results are (mainly) based on data for that specific year; non-marked data points are (mainly) interpolations or extrapolations.*

sumption remains constant to obtain the per-device power consumption values for office desktops and laptops.

Based on a study on the carbon emissions associated with ICT in Australia [29] and the previously mentioned study on the energy consumption of consumer electronics in USA homes [19] we obtained an average annual energy consumption value for CRTs and LCDs.

2.3.5 Result

The final results of our calculations are given in Table 2.2 and shown in Fig. 2.3. The total energy consumption by personal computers and their displays is currently around 300 TWh per year. The annual growth rate of this total electricity consumption over the time frame 2006 to 2012 is 5.3%. This growth rate is significantly lower than for device numbers (which is around 11-12% for computers and monitors), mainly due to the growing popularity of laptops and LCD monitors, which are more energy efficient than desktops and CRT monitors.

2.3.6 Reliability

In our calculations we had to make some assumptions where we could not find the required data. To the best of our knowledge, there are no recent statistics available for the time an average computer spends in active, sleep and off mode. These parameters have a considerable influence on the average power consumption of PCs[4], so it would be interesting to have more (worldwide) data available on this topic. Comparing the active, sleep and off times to the time PCs are actually used would also enable an estimation of the power savings that can still be achieved through the introduction of more intelligent power management. There are some estimates of the average energy consumption by PCs available in literature, but these are often national averages and are only available for more developed economies. This raises questions on the representativity of these values when we estimate the worldwide energy consumption. Furthermore, the values we did find in literature sometimes show a large spread. For example, according to [30], an average laptop in Europe consumed 116 kWh/y in 2007 and an average laptop in Switzerland consumed 47.5 kWh/y in 2008, while in [19] the average energy consumption of a laptop in the USA is estimated at 72 kWh/y. It is clear that further research in this area could greatly increase the reliability of our estimates.

While we have a reliable estimate for the number of personal computers in use based on UN statistics, the number of (external) computer monitors was harder to assess. Our calculations were complicated by the fact that shipping data for computer monitors is not publicly available. Our estimate for the number of devices is based on USA data solely, and would be more reliable if averages for different regions could also be found. Moreover, the availability of detailed shipping numbers would allow us to obtain more reliable estimates of the average power consumption of computer screens.

The influence of changes in individual parameters on the combined end result is however limited in most cases, which leads us to believe our results are a good indication of the worldwide annual energy consumption by personal computers and computer monitors.

2.4 Data Centers

In the data center category we cover the worldwide power consumption associated with computer servers, whether located in large data centers or

---

[4] For example, a PC that is fully on for 8 hours/day and 5 days/week (i.e., a typical work week) and turned off otherwise, consumes only about 25% of the power consumption of a PC that is on all the time.

Table 2.3: *Worldwide power consumption of data centers in 2007 and 2012. We adapted data from [31] by including orphaned servers, assuming no growth in the power per device since 2005, and assuming the server installed base grew from 2010-2012 as it did in 2005-2010. (VS = Volume server, MRS = Mid-range server, HES = High-end server)*

	2007			2012		
Server class	VS	MRS	HES	VS	MRS	HES
Power/server	222 W	607 W	8 106 W	222 W	607 W	8 106 W
Installed base	26.65 M	1.17 M	0.08 M	35.44 M	0.89 M	0.15 M
No of servers	Installed base × 1.25					
Storage pow.	24% of total server power consumption					
Comm. pow.	15% of total server power consumption					
PUE	1.95	1.95	1.95	1.83	1.83	1.83
Total power	176 TWh	21 TWh	19 TWh	219 TWh	15 TWh	34 TWh

in smaller spaces such as office server rooms. To estimate the total electricity used by data centers worldwide in the time frame 2007 to 2012, we base ourselves on the latest study by Koomey on this topic [31]. Koomey provides an estimation of data center power consumption for 2005, and a lower and upper bound estimation for 2010. We use newer data to estimate a most likely value instead of an upper and lower bound for 2010, and extend these trends to 2012. A key difference is that we include electricity use attributed to so-called ‘orphaned servers’, i.e., a typically undocumented number of servers using electricity but no longer delivering services.

The data center power consumption calculation follows the methodology outlined in [31]. To get the worldwide electricity consumption of *servers* we multiply, for each of three server classes, the average power per server by the number of servers worldwide. We then add the electricity used by *storage equipment* (tapes and hard disks), *communication equipment* (such as network switches) and *infrastructure equipment* (such as cooling and power provisioning losses) by applying three overhead factors. See also Table 2.3.

We consider Koomey’s (i.e., IDC’s) three cost-based classes of servers: volume servers (< $25 000 per unit), mid-range servers (between $25 000 and $500 000 per unit) and high-end servers (> $500 000 per unit). As the server count for these classes is based on commercial estimates, it does not account for custom-made servers from companies like Google or Amazon. Koomey has shown the impact of these servers to be still relatively small [31]. Custom-made servers might become a factor to consider in the future, however.

2.4.1 Electricity use per server

In [31], the 2010 lower bound scenario assumes no growth in power per server since 2005 (to reflect the industry's increased focus on energy-efficiency), whereas the upper bound scenario assumes the power per server trends from 2000 to 2005 extend beyond 2005.

To get recent data on the electricity use per server we used data available at spec.org, a non-profit corporation that establishes performance and power consumption benchmarks for computers. We analyzed the server power consumption (at 50% average target load) for all servers up to 1000 W in the spec.org power database [32] between January 2008 and December 2012 (393 entries). We created a volume and mid-range cluster by separating at 350 W (based on the power per server in 2005). We assumed 1000 W as an upper bound for the mid-range servers; few data points higher than this value were available anyway. The volume servers cluster (340 entries) shows a -3% Compound Annual Growth Rate (CAGR) in power per server in the period 2008-2012, and the mid-range servers cluster (53 entries) shows a 0% CAGR (i.e., no change) for the same period. The high-end cluster is not captured at all by the sample. A sensitivity analysis of the resulting volume and mid-range CAGR values to the cluster separation value, shows that the volume server CAGR is relatively stable at -3% (ranging from -3% to -2% in the cluster separation interval of 225 W to 475 W), whereas for the mid-range servers, the CAGR varies from an significant increase in power per server (6% per year, at 250 W cluster separation) to a negligible decrease.

We chose *not* to apply these CAGR values directly in our calculations since the spec.org sample is probably biased towards more energy-efficient servers. However, as volume servers dominate by far the server power consumption, these trends do suggest that the increase in power per server from 2000-2005 reported in [31] has not continued. Therefore we assume for all years in the time frame 2005 to 2012 the same power per server values as reported for the year 2005 in [31]. These values for each of the three server classes are listed in Table 2.3.

2.4.2 Worldwide number of servers

The worldwide number of servers for 2005 and 2010 is reported in [31]. For the worldwide number of servers in 2011 and 2012, we assume that the 2005 to 2010 server growth trends reported in [31] have continued to 2012. These trends showed a slower growth of volume servers (5.9% p.a.), a decrease in mid-range servers (-5.3% p.a.), and an increased growth of high-end servers (13.1% p.a.).

We assume continued trends based on the IDC server shipment data reported for 2011 [33] and 2012 [34]. The data we have available from IDC only details the total server shipments (i.e, the sum of all three server classes). However, the strong domination of the volume servers in the total number of shipped servers (for 2010, volume server shipments represented 98% of the total server shipments [31]) allows us to use the IDC data as indicative for volume server trends. The IDC data suggests a growth in the server installed base from 2010 to 2012 that is only slightly higher than the 5.9% p.a. rate observed from 2005 to 2010.

We adjust the number of servers above (i.e., the 'installed base') upwards with a factor 1.25 to account for orphaned servers, i.e., about 20% of the servers in many data centers are using electricity but no longer delivering computing services. In [31], orphaned servers are estimated to be 10-30% of the servers based on anecdotal evidence. Assuming an average value of 20%, this results in a factor of $20/80 = 25\%$ relative to the reported installed base.

Both the server worldwide installed base and the orphaned correction factor are shown in Table 2.3.

2.4.3 Storage, communications and infrastructure overhead

In line with [31], the storage and communication equipment power consumption is added as a fixed percentage of the server power consumption, i.e. 24% and 15% respectively.

The infrastructure equipment comprises cooling, power provisioning and power backup systems. Its power consumption is commonly captured by the Power Usage Effectiveness (PUE), a factor ≥ 1. For example, a PUE of 2 implies that for each watt of IT electricity use (i.e., by servers, storage and communication equipment), an additional watt is consumed by the infrastructure equipment. Koomey distinguishes in [31] an upper bound value of 1.92 (based on [35]) and a lower bound value of 1.83 (based on [36]). We assume an average PUE of 1.88 for the year 2010. Based on a PUE of 2 for the year 2005, we linearly interpolated the intermediate years, and linearly extrapolated this trend for the years beyond 2010. This results in a PUE of 1.95 for 2007 and a PUE of 1.83 for 2012, as shown in Table 2.3.

2.4.4 Result

Our results show that data centers worldwide consume 270 TWh in 2012, as shown in Fig. 2.4. The CAGR from 2007 to 2012 is 4.4%. The data center power consumption is dominated by infrastructure electricity use (i.e., cooling and power supply losses). The actual server power consumption

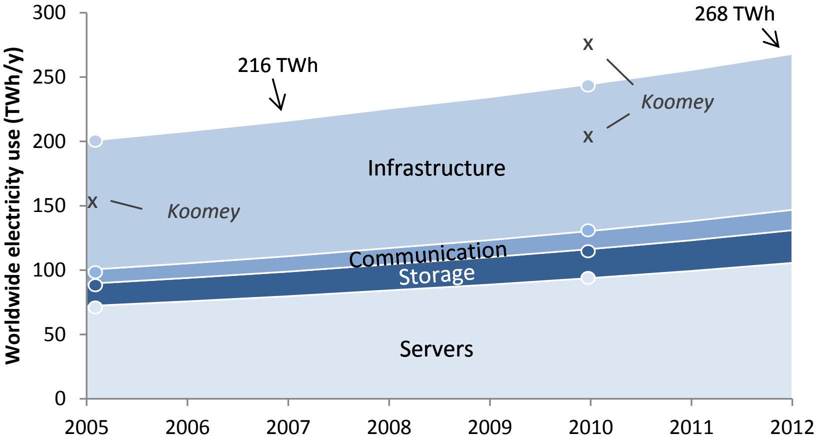

Figure 2.4: *Worldwide use phase electricity consumption of data centers. Infrastructure electricity use (mainly cooling and power supply losses) dominate the result, followed by the electricity used by the actual servers. The three crosses indicate Koomey's values for 2005 and 2010 (upper bound an lower bound) [31]. While our results in general follow Koomey's lower bound assumptions, they are shifted upwards (and, incidentally, close to the average of the 2010 lower and upper bound value) since we account for orphaned servers.The circular markers ○ indicate years for which the subcategory results are (mainly) based on data for that specific year; non-marked data points are (mainly) interpolations or extrapolations.*

accounts for only about 40% (see Fig. 2.4), and is clearly dominated by the share of volume servers (see Table 2.3).

2.4.5 Reliability

Our estimation is mainly based on [31], as it provides the most substantiated values available on this subject. There are few studies that provide an estimate for the worldwide data center electricity use, and most of them either base themselves on the same study (or an earlier publication from the same author), or are outdated for the time frame we consider.

An important uncertainty in [31] is related to the power per server in 2010. Our analysis of spec.org data leads us to believe that a stagnation in power per server since 2005 is more likely than a continuation of the pre-2005 power-per-server trends.

The PUE value is a second factor that is rather uncertain. As any changes to the PUE apply linearly to the result, the impact of any deviation is potentially large. A worldwide data center survey conducted in

2012 by the UptimeInstitute [37] reports that the PUE reported by its participants averages between 1.8 and 1.89. While this value might be biased (as data centers with a focus on energy-efficiency are more likely to report their results) it is in line with the value we extrapolated for 2012, i.e. 1.83, which increases the confidence in our results.

Finally, our accounting for orphaned servers based on a value that was reported, but not used, by Koomey, might raise some criticism, and rightly so. However, we think the *actual* worldwide power consumption is better approximated by including it. As it is applied as a single factor across all years considered, it does not influence the observed growth trend.

2.5 Overall trends and observations

In the previous sections we discussed the electricity use of three categories—communication networks, personal computers, and data centers—separately. Here, we compare their trends and absolute power consumption values to each other, to remaining ICT equipment categories (such as TVs and mobile phones) and to the total worldwide electricity use.

All values in this section apply to the time frame 2007 to 2012.

2.5.1 Growth trends

Communication networks show the highest increase of electricity use, with a CAGR of 10.4% (see Fig. 2.5). The growth rates of PCs and data centers are both only about half of that value. All three growth rates are higher than the growth rate of the total worldwide electricity consumption (about 3% per year[5]) [38]. This implies that the share of these ICT categories in the total worldwide electricity consumption is increasing year after year.

The observed growth rates are lower than what we projected in our earlier study by Pickavet et al. [5] in 2008; we then estimated the growth in a business as usual scenario to be 12% per year for communication networks, nearly 8% per year for personal computers, and 12% per year for data centers. While part of this difference might be attributed to the uncertainty associated with our estimates, we see two other potential reasons for this significant decrease in growth rates. First, the increased attention for more energy-efficient technologies has brought down the electricity use growth rates. This is clearly visible in the personal computer category, with the shift

---

[5] The CAGR of the worldwide electricity consumption from 2000 to 2011 is 3.4%. We used this long-term trend to extrapolate the 2011 value to 2012. The CAGR from 2007 to 2012 is slightly lower at 2.9% because of the impact of the global financial crisis in 2008 and in the subsequent years.

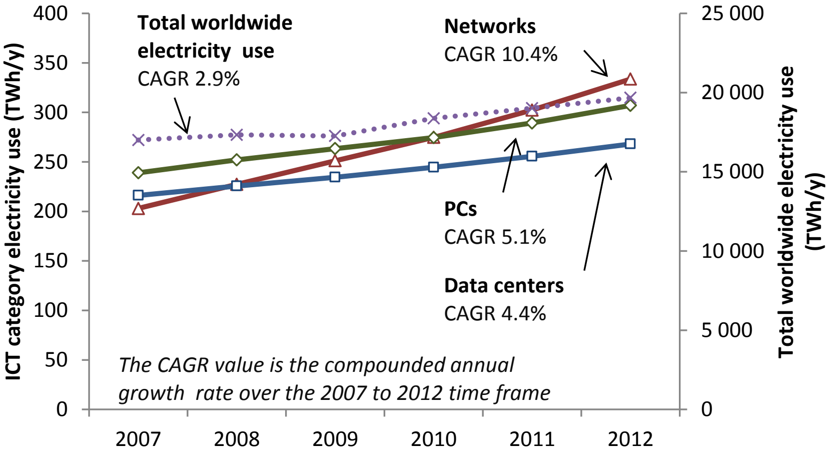

Figure 2.5: *Evolution of worldwide electricity use of networks, PCs and data centers (solid lines, left axis) and total worldwide electricity use (dotted line, right axis). Over the last five years, the electricity consumption in all three ICT categories increased at a rate higher than the total worldwide electricity consumption. In 2012 each category accounts for roughly 1.5% of the worldwide electricity consumption. Note that, since some of the data points between 2007 and 2012 are based on interpolations, small variations might not show up in the intermediate years.*

from CRT to LCD monitors, and from desktop to laptop computers. Second, it is not unlikely that the global financial crisis from 2008 had an impact on the buying behavior of end-users and businesses related to ICT equipment, and consequently the associated electricity use. It is important to point out that—as the main intention of our study was to capture the growth trend from 2007 to 2012—some data points are the result of interpolation and extrapolation. Therefore, variations in intermediate years might not show up in our results.

Communication networks We cannot easily attribute the relatively high growth in communication networks to a specific factor. Telecom operator networks dominate the communication network power consumption (see Fig. 2.2), and they are the main driver for the growth. However, as explained earlier in Section 2.2.1, our methodology does not allow for a further breakdown of power consumption across mobile, fixed broadband and fixed telephony services. The worldwide number of mobile subscriptions and fixed broadband subscriptions show an increase of 13% to 14% per year [7], the number of fixed telephony subscribers has decreased at a rate of 1% per year. With the growth rate of mobile and fixed broadband subscriptions being somewhat higher than the growth rate of the electricity use in telecom operator networks, it is likely that the electricity consumption per average subscriber is decreasing. It is however not clear if this can be attributed to the (intentional) replacement of old equipment with more energy-efficient devices, or if this is rather because of a shift to new technologies such as mobile communication.

Personal computers The relatively modest growth in PC power consumption is attributed to a shift to more power-efficient technologies, notably from desktops to laptops and from CRT monitors to LCDs. For this reason, while the number of computers and monitors have grown at a rate of 11-12%, total PC electricity use has grown at a rate of just over 5% per year.

Data centers The increase in the number of servers drives the growth in data center electricity use, despite the slight improvement in PUE. As volume servers dominate the data center power consumption (see Table 2.3), their growth rate (at 5.9% p.a.) drives the overall data center electricity growth rate. The growth rate for data centers (5.1% per year) is significantly lower than what we estimated five years ago (12% per year). This is due to a reduced growth in number of servers (caused by the 2008 financial crisis, the associated economic slowdown, and further improvements in virtualization [31]) and the assumed stagnation in power per server. Especially in medium to large data centers, the incentive for actions to improve the energy-efficiency can lead to very visible reductions in electricity use (and associated costs), while these actions are at the same time relatively easy to implement due to economies of scale. In locations with only a few servers, on the other hand, the server electricity use is probably a relatively minor cost, and consequently less of a focus for improvements or optimizations.

Cross-domain observations In [39] it is observed that 'the electrical efficiency of computation has doubled roughly every year and a half over the last six decades'; this corresponds to an annual reduction of 38% in power per unit of computation. It is interesting to note that this trend—which applies to laptops, desktops and servers in our study—has not resulted in a reduction or even status quo of the overall power consumption of these devices. Instead, it appears that precisely this power-efficiency trend leads to the emergence of new technologies such as laptops and mobile phones. These, in turn, have led to new services and applications leading to an overall increase in ICT power consumption. In this light, it is interesting to observe that our results do not indicate (yet) that laptops have replaced desktops (in Fig. 2.3 there is no decline in desktop power consumption);

laptops appear to be used *in addition* to the already existing worldwide user base of desktops.

Similarly, anecdotal evidence suggests that network operators have a tendency to build new networks on top of existing networks, leaving older equipment in place for supporting legacy devices and services. This suggests that the equipment lifetime for network equipment is much longer than for PCs and servers, which might partly explain the higher growth rate for power consumption in communication networks when compared to the other two categories.

2.5.2 Relative power consumption

The three categories were each consuming roughly an equal amount of power in 2007. The relatively high growth rate of communication networks electricity use has lead to this category to overtake the power consumption of both PCs and data centers in 2012 (see Fig. 2.5).

Still, each of these categories only accounts for a small share of the total worldwide electricity consumption, respectively 1.7% for networks, 1.6% for PCs and 1.4% for data centers in 2012.

2.5.3 Power consumption of the remaining ICT equipment

For reference, and to provide a 'bigger picture' view, we have also tried to estimate the power consumption of the remaining ICT equipment. We explicitly point out that the estimates below are provided to give a rough indication only. We have grouped these in two categories, i.e. TVs and Others.

TVs We estimate the worldwide TV power consumption in 2012 to be in the order of 400 – 500 TWh. This estimation is mainly based on combining the results in [19] (which provides a detailed estimation of TV electricity use in the USA) and [40] (which estimates the worldwide TV electricity consumption based on present and future TV energy-efficiency levels, but doesn't seem to take into account the electricity consumption of legacy TVs).

Other ICT equipment We estimate the worldwide power consumption of other ICT equipment to be in the order of 300 – 500 TWh. While the 'others' category by definition comprises all remaining ICT equipment, we have based ourselves on the OECD definition of ICT[6], but included in our estimation only those categories which we believe to represent the bulk of the electricity consumption. More specifically we have included the following equipment in our estimate (roughly from largest to smallest share): set-top boxes, general other household ICT equipment (such as radios, hi-fi systems, cordless phones, docking stations and VCRs), DVD and blu-ray players/recorders, video game consoles, general office ICT equipment (such as printers, faxes, scanners, telephony equipment and audio/visual equipment), mobile phones, computer speakers, household printers, and ATMs. Note that we did not take into account point-of-sale terminals, PDAs, burglar or fire alarms, TV cameras, tablets and electronic integrated circuits integrated in devices generally not considered as ICT (for example, the control electronics inside washing machines). Tablets are not specifically identified yet in the OECD definition (although they fit under 'portable automatic data processing machines weighing not more than 10 kg, such as laptop and notebook computers'); however, we show in Section 2.6.2 that their electricity use is still negligible.

Our rough estimates for TVs and Other ICT equipment suggest that our three categories—communication networks, personal computers and data centers—account for about half of the worldwide ICT electricity use. This would mean that the use phase of ICT accounts for around 9% of the total worldwide electricity consumption in 2012. However, as we mentioned earlier, the above statements have to be treated with caution, since our estimates for TVs and Other ICT equipment are only rough estimations to perform a first-order comparison.

2.6 Comparison with other studies

In this section we compare our results to those of a number of other studies in order to get an indication of the validity of our results. We look at an earlier estimate of our research group [5] published in 2008, the SMARTer2020 report published in 2012 [6] which is a follow-up to the well-known SMART2020 report from 2008 [3], and work published in 2010 by Malmodin et al. [2].

Fig. 2.6 shows our results for the year 2012 and the results of the four said works for their respective applicable year. The bars represent the

---

[6] As ICT starts to entangle all aspects of our lives, the scope of what exactly ICT equipment (and services) comprises is becoming increasingly difficult to define. The OECD probably provides the best documented scope of 'ICT products' in Table 2.A1.1 of [41]. Another approach, proposed and used in [42], would be to abandon the usage of the term ICT altogether, and consider instead 'Electronics' which is defined as 'any device whose primary function is information'.

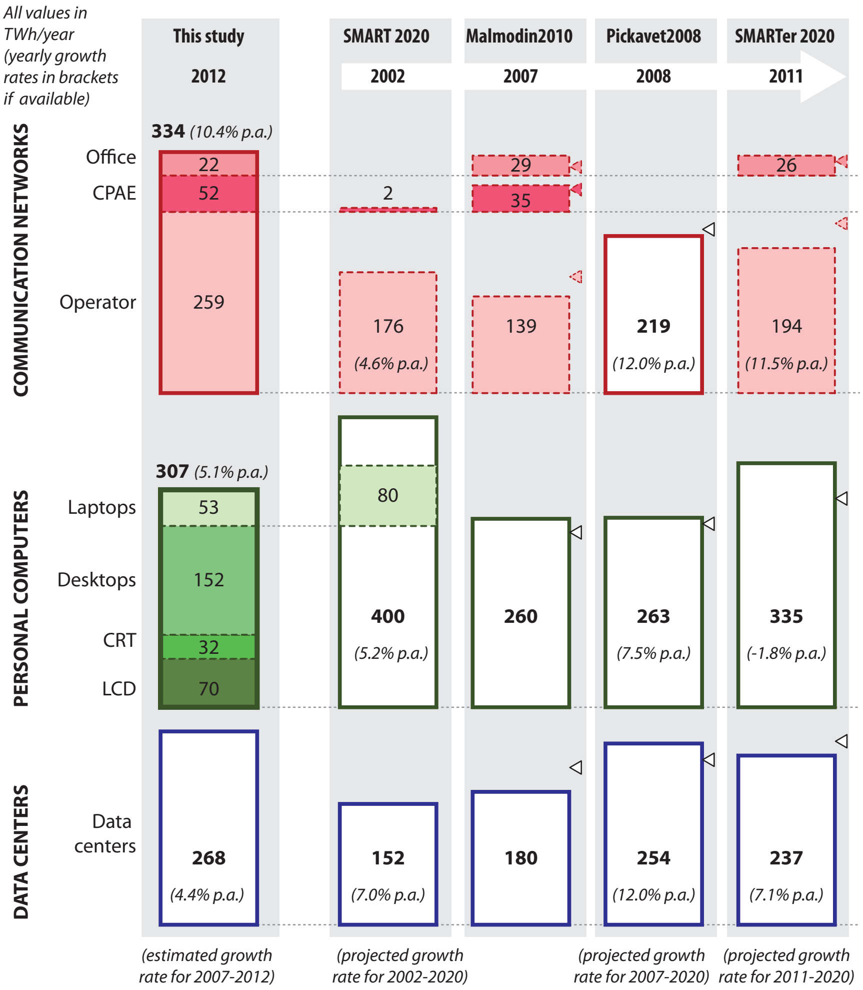

Figure 2.6: *A comparison of our 2012 results with a number of other often-cited studies [2, 3, 5, 6]. These works have been ordered by year of applicability. The triangles to the right of the boxes indicate, for easier comparison, our estimation for the applicable year. The SMART2020 report values have been derived from the use phase CO_2 values assuming 500 gCO_2/kWh.*

Table 2.4: *Comparison of historical and projected annual growth rates in worldwide electricity use. Our estimates for 2007 to 2012 are lower than all projections for 2020, two exceptions notwithstanding.*

	This study *(estimated 2007-2012)*	SMART2020 *(projected 2002-2020)*	Pickavet2008 *(projected 2007-2020)*	SMARTer2020 *(projected 2011-2020)*
Networks	10.4%	4.6% [a]	12.0%	11.5% [a]
PCs	5.1%	5.2%	7.5%	-1.8%
Data centers	4.4%	7.0%	12.0%	7.1%

[a] Telecom operator network subcategory only

electricity use (in TWh) for each of the three categories, broken down into subcategories if detailed values were available or could be derived. The observed or projected annual growth rate is indicated between brackets. Triangles indicate our estimate for the corresponding year. For clarity, the observed and projected annual growth rates are also reported in Table 2.4.

The works by Raghavan [43] and Somavat [44] are two recent studies that provide estimates for the ICT categories we consider. Although we mention them here for completeness, the estimates in [43] are too crude for our purpose. The estimates in [44] are roughly in line with our findings, which is not surprising as they are partly based on our earlier work [5] and Koomey's work [45].

2.6.1 Communication networks

Note: a more detailed comparison of communication networks electricity use is available in [7]. It contains a wider set of related works, but does not include the SMARTer2020 report [6] which was not yet available at the time.

It is clear from Fig. 2.6 that our estimate of network electricity use is (significantly) higher than earlier/other estimates. We think the main reason for this is because our methodology captures the (hidden) overhead associated to operators in a more accurate way.

The 2007 value from Malmodin et al. [2] for the CPAE subcategory is in line with our results for 2007, being 31 TWh. Their value for office networks is more than twice as high as our estimation (i.e., 12 TWh). Malmodin assumes a fixed overhead of 8 watts for each office PC, which—as we mentioned in Section 2.2.5—potentially includes office end-user equipment (such as faxes) that is out of our scope. The 2007 value from Malmodin for operator networks (139 TWh) includes overhead for offices and stores. To bring this in agreement with our scope, we should subtract this overhead, which we estimated in [7] to be about 13%. If we do so, their result for operator networks (121 TWh) is about 25% lower than our value (i.e., 160 TWh for 2007). This difference can probably be attributed to the fact that they used a different sample and did not distinguish between fixed broadband and fixed telephony users in their calculation method.

The 2008 value from our earlier work (Pickavet et al. [5]) is only about 2% lower than our current value for 2008. On the other hand, because of the adjusted growth rates, our earlier estimate for 2012 is 3% higher than our current value for 2012.

As the SMARTer2020 report [6] based its network electricity consumption estimation on the study by Malmodin, we see similar deviations from our results, i.e. a lower estimation for the operator subcategory, and a higher estimation for the office subcategory[7]. However, the deviation in the office subcategory from our results is only 30%, which is lower than what we observed for Malmodin's study. This is because the SMARTer2020 report assumed 4 watts for network equipment per office PC, or half of Malmodin's value[8]. While 4 W/unit is slightly lower than our 4.4 W/unit for 2011, the higher number of office PCs in the SMARTer2020 report results in an overall higher value for office network electricity consumption.

2.6.2 Personal computers

We have a consistently lower estimate for the worldwide electricity use by personal computers and monitors compared to the other works. The main reason seems to be that we estimate the average yearly electricity use per PC setup lower than other works. An overview is given in Table 2.5. This is somewhat surprising, as we are using USA-based data, which we thought might lead to an overestimation. A study dedicated to mapping variations in the average yearly PC electricity consumption across different worldwide regions would certainly be beneficial to clear up this issue.

Malmodin's [2] value for the worldwide electricity use of PCs in 2007 is about 9% higher than our value. While the worldwide total number of computers for 2007 is very similar, Malmodin's average power consumption per PC (including monitors) is approximately 13% higher than our average value for the same year.

The estimate in our earlier work [5] for 2008 is 4% higher than our

---

[7] We do not have a SMARTer2020 report value for the CPAE subcategory. The SMARTer2020 report states that 'set top boxes, home routers and modems and other computer peripherals' account for 13% of the overall end-user device emissions. This cannot be mapped directly to electricity use of our CPAE subcategory.

[8] This is probably because the SMARTer2020 report left out office end-user equipment such as faxes, which make up about half of Malmodin's 8 W/unit.

Table 2.5: *Comparison of the average energy consumption (kWh/year) per PC (inc. monitor). For this study, the values are obtained by dividing the total power consumption of laptops, desktops and monitors by the total number of laptops and desktops.*

Applicable year	2007	2008	2011	2012	2020
This study	221	208	177	169	-
Malmodin2010 [2]	250	-	-	-	-
Pickavet2008 [5]	-	263	-	-	-
SMARTer2020 [6]	-	-	219	-	102

current value for 2008. The shift to more power-efficient technologies such as laptops and LCD monitors explains the reduced annual growth rate in electricity use; our earlier work was based on the number of computers growing at about 10% per year, which is similar to our current observed value of 11-12% (see Section 2.3.5).

The 2011 estimate by the SMARTer2020 report [6] is 16% higher than our estimate. While their assumed installed base of PCs is slightly lower than our numbers, again the higher electricity use per personal computer results in a higher total worldwide electricity use for personal computers. Interestingly, they forecast a -1.8% compounded annual decrease for PC electricity use from 2011 to 2020[9]. The reason behind this downward trend is the halving of the average electricity use per device by 2020, driven by both 'efficiency gains and fewer hours spent on PCs due to the emergence of smart devices, and a greater use of laptops vs. desktops' [6]. We can not yet observe this trend in our estimates, see Fig. 2.3.

With the recent explosion of tablet device sales—the first Apple iPad was released mid 2010, and over 100 million devices had been sold by the end of 2012 [46]—it might be interesting to point out that the SMARTer2020 report estimates the tablet worldwide electricity use in 2011 at 1.1 TWh (but sharply increasing at 36% annually towards the year 2020). Even at a projected 1.5 TWh in 2012, this is—for now—still a negligible 0.5% fraction of the total electricity use of the PCs category.

2.6.3 Data centers

It is important to point out that all works depicted in Fig. 2.6 (including this study) based their data center electricity use estimation on work by

---

[9] In contrast, the *total* CO_2 emissions by PCs, consisting of both the electricity usage carbon emissions and the embodied emissions, is forecast by the SMARTer2020 report to increase with 1.2% per year from 2011 to 2020. This is due to the embodied emissions in the increasing number of shipments. It is not clear whether shipments are estimated to increase due to a growing user base, shorter device lifetimes, or both.

Koomey et al. [47] [45] [31]. There is good reason for this, as his work is backed up by solid data (that might otherwise be very hard to have access to) and a transparent methodology. With this in mind, we would expect consistent results across all works. However, we can see in Fig. 2.6 that this is not always the case. The reason is twofold.

First, in this study we have accounted for the electricity use of 'orphaned servers', i.e., a typically undocumented number of servers using electricity but no longer delivering services. While Koomey himself provides an estimate of this share, he considers the value too unreliable to include in his calculations. Nonetheless, we think the actual worldwide electricity use of data centers is better reflected when we include this share, which represents an additional 25% in electricity use over the non-inclusion scenario. This explains the lower estimates for 2008 by Malmodin [2] and for 2011 by the SMARTer2020 report [6]. The difference with the SMARTer2020 report value (237 TWh) is less than the expected 25%; the reason is that the report uses the average of Koomey's upper and lower bound scenarios, while our results are based on the lower bound power-per-server values.

Second, our earlier work (Pickavet et al. [5]) was based on initial work by Koomey in [47] which provided an estimate for the year 2005. Later on however, Koomey refined his estimate for 2005 in [45] and [31], especially with respect to the estimation of the storage and communications overhead. The result was that we overestimated this overhead for 2008. In addition, Koomey showed in [31] that the growth in electricity use after 2005 was not as high as projected earlier, which again lead to an overestimation for 2008 from our side. However, the overestimation in our earlier work compared to this current study is not as high as could be expected from both these hindsight observations, because it is dampened by the effect of accounting for orphaned servers, as described above.

2.7 Conclusion and outlook

2.7.1 Conclusion

Growth trends The combined electricity consumption of communication networks, personal computers and data centers is growing at a rate of nearly 7% per year (i.e., doubling every 10 years). The strongest growth is observed in communication networks, at 10% per year, probably fueled by the increase in (mobile) interconnectivity of digital equipment. The electricity consumption of personal computers is growing at 5% per year, and that of data centers at 4% per year. All growth rates have decreased compared to what we predicted in a similar study five years ago. This can partly be attributed to a shift to more energy-efficient technologies (such as from CRT to LCD monitors, and the introduction of server virtualization), and potentially to the effects of the global financial crisis in 2008.

Absolute power consumption Together these three ICT categories consumed about 900 TWh in 2012. The relative share of these ICT products and services in the total worldwide electricity consumption has increased from about 3.9% in 2007 to 4.6% in 2012. This does not yet include the electricity consumption of other devices that are usually considered as part of ICT, such as TVs and their set-top boxes, (smart) phones and audio devices.

Focus of reduction efforts The electricity consumption of each of the three considered categories is about the same size. This highlights the need for energy-efficiency research across all three domains, rather than focusing on a single one. On the other hand, it might be useful for future work to rethink the breakdown of ICT electricity use in the presently considered categories. Perhaps an assessment on e.g. displays or general overhead (such as power-supply and standby losses), might lead to new insights on where the main focus should be.

Comparison to other studies Our results are also consistent with other research on this topic. A notable exception is that we consistently estimate the electricity use attributed to telecom operator networks higher than other works. We attribute this to our methodology which we think is more inclusive and representative of the actual electricity use in this subcategory.

2.7.2 Reflections and outlook beyond 2012

While the last decade has seen an increasing attention for energy-efficiency, this has not yet translated in an absolute reduction or even status quo of the total ICT electricity use. Indeed, the growth of some specific (sub)categories has slowed down, but there is a shift to new applications and technologies such as LCDs, laptops, and tablets. While these devices have smaller electricity usage, this energy-efficiency improvement is outweighed (or soon could be) by a fast growth in device numbers. In this light, it might be interesting to research to which extent an increase in energy-efficiency has made the continued growth in ICT services possible, and thus partly fueled the associated growth in ICT electricity use as well.

Finally, looking beyond 2012, it is difficult to predict future growth rates. Personal computers will probably become even more efficient as the world continues to move to more mobile forms of end-user computing devices, potentially resulting in a stabilization or even decline in total electricity use. The electricity consumption in communication networks might continue to increase at the current rate, with more and more devices being connected and the advent of machine-to-machine communication (for example, your electric utility company could be telling your washing machine to start when a surplus of renewable energy becomes available). In addition, while the percentage of individuals using the Internet in emerging economies such as China and India is steadily rising, it is still far below those in more mature markets such as the USA, Western Europe and Japan. Concerning data center power consumption, we see two opposite trends. The increasing popularity of cloud-based computing and storage could result in more servers (and an associated increase in electricity consumption). On the other hand, this might also be an opportunity to move servers running at low overall efficiency in small offices and companies to more energy-optimized data centers.

In the future, frequent estimates of the worldwide electricity use by ICT will be essential to provide timely feedback if indeed ICT electricity consumption remains relatively small, or instead continues to grow at an unsustainable rate.

Acknowledgments

The work described in this paper was carried out with the support of the iMinds-project GreenICT, the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 247674 (STRONGEST project), n. 257740 (Network of Excellence TREND), and n. 288021 (Network of Excellence EINS). Sofie Lambert is funded by the Agency for Innovation by Science and Technology in Flanders (IWT).

References

[1] G. A. Fowler. *Facebook: One Billion and Counting*. The Wall Street Journal, October 4, 2012. Available from: http://online.wsj.com/article/SB10000872396390443635404578036164027386112.html.

[2] J. Malmodin, A. Moberg, D. Lundén, G. Finnveden, and N. Lövehagen. *Greenhouse Gas Emissions and Operational Electricity Use in the ICT and Entertainment & Media Sectors*. J. Ind. Ecol., 14(5):770–790, 2010. doi:10.1111/j.1530-9290.2010.00278.x.

[3] M. Webb et al. *SMART2020: Enabling the low carbon economy in the information age*. The Climate Group, 2008. Available from http://www.smart2020.org/_assets/files/02_Smart2020Report.pdf.

[4] P. Teehan and M. Kandlikar. *Sources of variation in life cycle assessments of desktop computers*. Journal of Industrial Ecology, 16(s1):S182–S194, 2012.

[5] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt, and P. Demeester. *Worldwide energy needs for ICT: the rise of power-aware networking*. In 2nd International Symposium on Advanced Networks and Telecommunication Systems, 2008 (ANTS '08), pages 1–3, December 2008. doi:10.1109/ANTS.2008.4937762.

[6] GeSI and BCG. *SMARTer2020: The Role of ICT in Driving a Sustainable Future*, 2012. Available from: http://gesi.org/SMARTer2020.

[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, December 2012. doi:10.1364/OE.20.00B513.

[8] D. Kilper, G. Atkinson, S. Korotky, S. Goyal, P. Vetter, D. Suvakovic, and O. Blume. *Power Trends in Communication Networks*. IEEE Journal of Selected Topics in Quantum Electronics, 17(2):275 –284, march-april 2011. doi:10.1109/JSTQE.2010.2074187.

[9] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker. *Energy Consumption in Optical IP Networks*. J. Lightwave Technol., 27(13):2391–2403, Jul 2009. doi:10.1109/JLT.2008.2010142.

[10] S. Lanzisera, B. Nordman, and R. Brown. *Data network equipment energy use and savings potential in buildings*. Energy Efficiency, 5:149–162, 2012. doi:10.1007/s12053-011-9136-4.

[11] HomePlug PowerLine Alliance. *HomePlug Certifies Products Based On Newly Ratified IEEE 1901 Powerline Networking Standard*, October 2010. Press release available from: https://www.homeplug.org/news/pr/view?item_key=871577ee4d3afd08e7d9ccf1a6300061add3681a.

[12] *dLAN 200 AVplus Starter Kit*, 2012. Data sheet available from: http://www.devolo.com/consumer/dlan-200-avplus/pdf/datasheet-dlan-200-avplus-starter-kit-en.pdf.

[13] Solwise. *NET-PL-200AV-MINI*, 2013. Available from: http://www.solwise.co.uk/net-powerline-av-mini.htm.

[14] *Global ICT developments, 2001-2011*. ITU World Telecommunication / ICT Indicators database, 2011. Available from: http://www.itu.int/ITU-D/ict/statistics/material/excel/20112/ictwebsite/Global_ICT_Dev_01-11.xls.

[15] *World population prospects: the 2010 revision*. UN Population Division, 2011. Available from: http://data.un.org/Data.aspx?q=population&d=PopDiv&f=variableID%3a12.

[16] *World broadband statistics: Q4 2007*. Point-Topic, 2008.

[17] *World broadband statistics: Q4 2011*. Point-Topic, 2012.

[18] *The little data book on information and communication technology*. The World Bank / ITU, 2007-2012.

[19] B. Urban, V. Tiefenbeck, and K. Roth. *Energy consumption of consumer electronics in U.S. homes in 2010.* Fraunhofer Center for Sustainable Energy Systems, December 2011.

[20] C. Cremer, W. Eichhammer, M. Friedewald, P. Georgieff, S. Rieth-Hoerst, B. Schlomann, P. Zoche, B. Aebischer, and A. Huser. *Energy consumption of information and communication technology (ICT) in Germany up to 2010: summary of the final report to the German federal ministry of economics and labour*, 2003.

[21] W. Vereecken, W. Van Heddeghem, M. Deruyck, B. Puype, B. Lannoo, W. Joseph, D. Colle, L. Martens, and P. Demeester. *Power consumption in telecommunication networks: overview and reduction strategies*. IEEE Commun. Mag., 49(6):62–69, June 2011. doi:10.1109/MCOM.2011.5783986.

[22] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch. *Energy consumption of telecommunication networks and related improvement options*. IEEE Journal of Selected Topics in Quantum Electronics, 17(2):285–295, 2011. doi:10.1109/JSTQE.2010.2053522.

[23] K. Kawamoto, J. Koomey, B. Nordman, R. Brown, M. Piette, M. Ting, and A. Meier. *Electricity used by office equipment and network equipment in the US*. Energy, 27(3):255–269, 2002. doi:10.1016/S0360-5442(01)00084-6.

[24] *Definitions of world telecommunication/ICT indicators*. ITU, March 2010. Available from: http://www.itu.int/ITU-D/ict/material/TelecomICT_Indicators_Definition_March2010_for_web.pdf.

[25] *Millennium development goals indicators*. UN Statistics Division, 2010. Available from: http://mdgs.un.org/unsd/mdg/Data.aspx.

[26] Asymco. *A brief history of personal computer vendors*. Available from: http://www.asymco.com/hire-me/historic-computer-shipments/, 2011.

[27] M. Murphy and M. Meeker. *KPCB top mobile internet trends*. Available from: http://www.slideshare.net/kleinerperkins/kpcb-top-10-mobile-trends-feb-2011, October 2011.

[28] F.-K. Wang and K.-K. Chang. *Modified diffusion model with multiple products using a hybrid GA approach*. Expert Systems with Applications, 36(10):12613 – 12620, 2009. doi:10.1016/j.eswa.2009.05.018.

[29] *Carbon and computers in Australia - the energy consumption and carbon footprint of ICT usage in Australia in 2010*. ACS, 2010.

[30] A. Andrae and O. Andersen. *Life cycle assessments of consumer electronics - are they consistent?* The International Journal of Life Cycle Assessment, 15(8):827–836, 2010. doi:10.1007/s11367-010-0206-1.

[31] J. G. Koomey. *Growth in data center electricity use 2005 to 2010*. Available from: http://www.analyticspress.com/datacenters.html, 2011.

[32] spec.org. *SPECpower ssj2008 benchmark*. http://www.spec.org/benchmarks.html, 2013. Data used up to (and including) December 2012.

[33] IDC. *Despite a 7.2% Decline in Fourth Quarter Revenue, Worldwide Server Market Revenues Increase 5.8% in 2011, According to IDC*, February 2012. Press release available from: http://www.idc.com/getdoc.jsp?containerId=prUS23347812.

[34] IDC. *Worldwide Server Market Rebounds Sharply in Fourth Quarter as Demand for x86 Servers and High-end Systems Leads the Way, According to IDC*, February 2013. Press release available from: http://www.idc.com/getdoc.jsp?containerId=prUS23974913.

[35] *Performance Ratings Technical Methodology for Data Center*. Energy Star, June 2010. Available from: http://www.energystar.gov/ia/business/evaluate_performance/data_center_tech_desc.pdf.

[36] *2011 Data Center Industry Survey*. Uptime Institute, 2011.

[37] M. Stansberry and J. Kudritzki. *Uptime Institute 2012 Data Center Industry Survey*. Uptime Institute, 2012. Available from: http://uptimeinstitute.com/2012-survey-results.

[38] *Global energy statistical yearbook 2012*. Enerdata, 2012. Available from: http://yearbook.enerdata.net/world-electricity-production-map-graph-and-data.html.

[39] J. G. Koomey, S. Berard, M. Sanchez, and H. Wong. *Implications of historical trends in the electrical efficiency of computing*. Annals of the History of Computing, IEEE, 33(3):46–54, 2011. doi:10.1109/MAHC.2010.28.

[40] W. Park, A. Phadke, N. Shah, and V. Letschert. *TV Energy Consumption Trends and Energy-Efficiency Improvement Options*. Technical Report LBNL-5024E, Ernest Orlando Lawrence Berkeley National Laboratory, July 2011. Available from: http://www.superefficient.org/en/Resources/~/media/Files/SEAD%20Televisions%20Technical%20Analysis.pdf.

[41] OECD. *OECD Guide to Measuring the Information Society 2011*. OECD Publishing, 2011. doi:10.1787/9789264113541-en.

[42] B. Nordman and S. Lanzisera. *Electronics and network energy use: Status and prospects*. In IEEE International Conference on Consumer Electronics (ICCE) 2011, pages 245–246, 2011. doi:10.1109/ICCE.2011.5722564.

[43] B. Raghavan and J. Ma. *The energy and emergy of the internet*. In Proceedings of the 10th ACM Workshop on Hot Topics in Networks, page 9. ACM, 2011. doi:10.1145/2070562.2070571.

[44] P. Somavat and V. Namboodiri. *Energy consumption of personal computing including portable communication devices*. Journal of Green Engineering, 1(4):447–475, 2011.

[45] J. G. Koomey. *Worldwide electricity used in data centers*. Environmental Research Letters, 3(3), July 2008. doi:10.1088/1748-9326/3/3/034008.

[46] B. X. Chen. *Apple, Facing Competition, Introduces a Smaller iPad of no significant change*. The New York Times, October 23, 2012. Available from: http://www.nytimes.com/2012/10/24/technology/apple-facing-competition-introduces-a-smaller-ipad.html.

[47] J. G. Koomey. *Estimating total power consumption by servers in the US and the world*, 2007. Available from: http://hightech.lbl.gov/documents/DATA_CENTERS/svrpwrusecompletefinal.pdf.

3

Power Consumption Modeling in Optical Multilayer Networks

In this chapter we lay the foundations for the work on backbone networks in the two subsequent chapters. We provide a set of representative power consumption values for backbone network equipment, which will be used in both Chapter 4 and Chapter 5. We also provide an analytical model to estimate the power consumption in a backbone network; this will be used for our quantitative survey of the power saving potential in backbone networks in Chapter 5.

⋆ ⋆ ⋆

Ward Van Heddeghem, Filip Idzikowski, Willem Vereecken, Didier Colle, Mario Pickavet, and Piet Demeester

Published in Photonic Network Communications, October 2012

Abstract The evaluation and reduction of energy consumption of backbone telecommunication networks has been a popular subject of academic research for the last decade. A critical parameter in these studies is the power consumption of the individual network devices. It appears that across different studies, a wide range of power values for similar equipment is used. This is a result of the scattered and limited availability of power values for optical multilayer network equipment. We propose reference power consumption values for Internet protocol/multiprotocol label switching (IP/MPLS), Ethernet, optical transport networking (OTN) and wavelength division multiplexing (WDM) equipment. In addition we present a simplified analytical power consumption model that can be used for large networks where simulation is computationally expensive or unfeasible. For illustration and evaluation purpose, we apply both calculation approaches to a case study, which includes an optical bypass scenario. Our results show that the analytical model approximates the simulation result to over 90% or higher, and that optical bypass potentially can save up to 50% of power over a non-bypass scenario.

3.1 Introduction

There is a growing number of publications on network power consumption It can be argued that interest and research into power consumption of Information and Communication Technology (ICT) networks started in 2003 with the paper 'Greening of the Internet' by Gupta and Singh [1]. At that time 'Green Networking' was still referred to as a 'somewhat controversial subject'. The paper discusses the power consumption of network devices and, on a larger scale, the Internet, and proposes a number of approaches to increase its energy-efficiency. Since then, numerous related papers have been published and presented. Most of these publications either provide an estimate of the current and future power consumption of (some subset of) networks, or evaluate a proposed solution for their power-saving potential. The main drivers for power reduction research are usually economical (reducing the energy cost), technical (reducing the associated heat dissipation) and environmental (reducing the carbon footprint) reasons.

Correct equipment power consumption values are key input for power evaluation studies All of the above purposes boil down to power consumption estimations, and one of the key inputs is the power consumption values of the constituting components. Sufficiently correct absolute power values are important for policy makers to assess the importance of ICT power consumption in comparison to other sectors. For example, if ICT networks consume relatively little power, it makes sense to focus research on using ICT networks to achieve energy savings in other domains. This is sometimes referred to as 'greening by ICT' and is the driver behind the frequently cited Smart 2020 report [2]. Sufficiently correct relative values of network equipment are important to network equipment vendors and researchers in order to focus on solutions with the largest overall saving potential. For example, as long as Optical Line Amplifiers (OLAs) constitute less than 3 percent of the total power consumption of a core network [3], there is little reason to focus research on making them more energy-efficient.

Currently used equipment power values suffer from a number of issues However, the power consumption values assumed in many papers suffer from a number of issues. First, they can differ substantially between publications. For example, while an optical amplifier is taken to consume 0.5 W per channel in [3] (the authors report 8 W per fiber, with a fiber carrying 16 channels), 1000 W per channel is assumed in [4]. This is more than three orders of magnitude difference. Second, one single device is often used as a source for the associated equipment power consumption, without being clear whether it is representative or not. In a few cases, no source is mentioned. Third, it is not always clear whether the power value used is just for the core functionality of the equipment, or whether it also takes into account any required control and support equipment like control cards and chassis power consumption. In addition, maximum power consumption values are sometimes used, which can differ substantially from power consumption under typical operating conditions.

To calculate total power consumption, simulation is not always practical The approach often used to estimate the total power consumption of a network with a given admissible topology fed with a certain traffic matrix, is based on dimensioning the network through simulation. Dimensioning entails determining the capacity requirements of all equipment. Simplifying the problem, dimensioning can be done by for example shortest-path routing all the traffic through the network. As a result of the dimensioning process all equipment counts (routers, router ports, transponders, etc.) are known. By multiplying the equipment count with the corresponding equipment power consumption the total network power consumption can be calculated. However, for large networks (in terms of nodes and links) this becomes computationally expensive. In addition, this approach does not give an indication upfront about the power consumption share of certain equipment and layers to the total result.

Contributions of this paper In this paper we address the issues outlined above for optical multilayer network. As such, the contributions of this paper are the following:

- we provide *reference values* for each equipment type, complete with direct source references where possible; the values are mostly based on public product data sheets (Section 3.3),

- we deduce a simplified *analytical power model* based on IP demands and the IP-layer hop count, that can be used as an alternative to dimensioning the network through simulation (Section 3.4),
- finally, in Section 3.5, we *illustrate and evaluate* with a case study how to use the information in this paper to determine the power consumption of an IP-over-WDM network, both via simulation and using the analytical model.

Due to space limitations, the individual reference values and detailed discussions are available as a separate report [5].

3.2 Related Work

We surveyed research articles that tackle cost models of multilayer networks. We looked at component-based and analytical power models, but considered also non-power consumption cost models.

Non-power consumption publications we build upon In [6], a Capital Expenditure (CapEx) model is given for optical multilayer networks, subdividing the network in four layers: Internet Protocol/Multiprotocol Label Switching (IP/MPLS), Ethernet, synchronous digital hierarchy/optical transport network (SDH/OTN) and Wavelength Division Multiplexing (WDM). Detailed normalized monetary cost values of equipment in each layer are listed in this paper. We use this model as a basis for our equipment categorization, updated to reflect recent changes and expected future evolutions. In [7], a so-called 'network global expectation model' is presented. The model proposes a number of equations to calculate expected values of network properties—such as the average node degree, the average number of hops, or the number of ports and capacity of a cross-connect—based on a few primary network properties. This approach is the idea behind the analytical power model we propose in Section 3.4.

Component-based power models Most of the publications evaluating solutions to increase energy-efficiency consider a power consumption model based on the individual power consumption of a few components and somehow counting the occurrence of each component (for example via a network dimensioning tool or Integer Linear Programming (ILP) approach). We provide a short selection of such publications here. In [3], the power saving possibility of static optical bypass over non-bypassed design in an IP over WDM network is investigated. The power consumption model considers IP router ports, transponders and optical amplifiers. In our related work on optical bypass [8], we assumed transponders to be part of the router interfaces, and additionally considered 3R regenerators. In [9], where optical cross connects are inserted between optoelectronic devices and the router in order to reduce power consumed in the network. Optical Cross-Connects (OXCs) and SONET/SDH devices are taken into account in addition to router ports, transponders and optical amplifiers. In [4], the energy-saving potential of turning off spare devices in an IP backbone network is investigated. The power model used is based on fixed-size core nodes with constant and equal power consumption and link power consumption (which is itself based on the inline amplifiers and the corresponding static power consumption of the router interface) scaling with the number of channels. In [10], Chabarek et al. measured the power consumption of two Cisco routers at different line card filling configurations. They devised a power consumption model from these observations that is the sum of the power consumption of the chassis and the installed active line cards (load dependent).

Analytical power models The following two works take a slightly different approach as they try to estimate the total power consumption rather than evaluate a specific solution for energy-efficiency. They calculate the total network power consumption directly, based on the average hop count and power efficiency values for the involved equipment. Additional factors account for traffic protection, future provision and cooling power overhead. In [11], Baliga, Tucker et al. propose a power consumption per customer model for optical networks, considering all main subnetworks such as access, metro and core. The power consumption in the core nodes is based on the power consumption efficiency of a typical core router. The link power consumption considers a channel efficiency value based on a typical WDM terminal system and inline amplifiers, differentiating between terrestrial and undersea links. In [12], a generalization of the model used in [11] is proposed, and referred to as a 'transaction-based model'. It is almost identical to the analytical power model we propose in Section 3.4, the main difference being that we consider a slightly different equipment breakdown and hop count attribution.

Other similar work The technical report by Idzikowski [13] provides an extensive list of power consumption values of various network elements of IP over WDM networks, based on product data sheets and research papers. The report categorizes the equipment in IP layer equipment and WDM layer equipment. The main difference with our work is that it does not homogenize the reported values based on for example functionality or capacity. In contrast, [14] uses a bottom-up approach to estimate the power consumption of high-capacity IP routers. It is based on aggregating the individual power consumption of the constituting parts such as transceivers, fabric interfaces and packet buffers. Different from our work, it is only focused on the nodes, rather than all network components. Power efficiency values are also given in [15], where a detailed analysis is done of various network element types (e.g., IP routers, Ethernet switches, SDH switches) and their functional components (framing, amplification, routing, etc.) with respect to power dissipation. However, in contrast to our work, it does not provide tractable references, and it does not include a power model.

3.3 Reference power consumption values

In this section we provide power consumption reference values for common IP over WDM equipment. These reference values are mostly based on publicly available product data sheets. Due to space restrictions, references to these source documents and associated detailed discussion for each equipment type are not given here. They are available in [5].

To provide consistent power consumption values, we provide:

- *typical values*, i.e., under typical load and conditions, rather than maximum power consumption values; please note that any derived efficiency values [W/Gbps] are calculated with respect to the capacity of the relevant equipment and not the actual throughput, which could be (far) less,
- values that *include chassis and control overhead power consumption*; external cooling or facilities overhead (lighting, etc.) is not included,
- values for *bidirectional equipment* (i.e., full-duplex).

Building on the CapEx work presented in [6], we consider the multilayer network and associated equipment to be subdivided in the following four layers:

- an *IP/MPLS layer* with associated routers which perform layer 3 switching,
- an *Ethernet layer*, which performs layer 2 switching,

- an *OTN layer*, which performs layer 1 time division multiplexing and transmission and adds monitoring,
- a *WDM layer*, which performs layer 1 space division multiplexing and transmission.

3.3.1 IP/MPLS layer

The IP/MPLS power consumption is based on publicly available data sheet values of two major commercial core routers: the Cisco CRS series and the Juniper T-series. More specifically, we will base the model on the values of the CRS-3 series, since it is the most recent architecture and most energy-efficient one (see Fig. 3.3).

Following the convention in [6], the equipment in the IP/MPLS layer consists of three building blocks (see Fig. 3.1). The *basic node* (e.g., a 1280 Gbps router) contains the chassis, switch fabric, routing engine, power supply, internal cooling and remaining minor components. The basic node contains *slot cards* (e.g., a 40 Gbps slot card), which contain one or more modules that can each hold a *port card* (e.g., a 4x10GE port card). The main functional block in the slot cards is the forwarding engine. The port card mainly contains the layer-2/3 interface and physical connection (such as PoS STM-256, or 10 Gigabit Ethernet).

This breakdown is representative for the power consumption of an IP/MPLS node. Fig. 3.2 shows the power distribution of five maximum core router configurations. The slot and port card combined make up roughly 75% of the power consumption. Power supply and internal cooling accounts for 10% (the CRS-3 value is lower because it does not include the power supply, which could not specifically be attributed to). Finally, the chassis is roughly 15%, mainly attributed to the switch fabric (about 10% of the total).

Table 3.1 lists the power consumption values for the various components, based on the CRS-3 router. The basic node building blocks consist of 16-slot line card shelves (LCSs) and optionally fabric card shelves (FCSs). The fabric card shelf can connect up to 9 line card shelves, and a configuration with maximum 8 fabric card shelves (and thus 72 line card shelves) is possible. The table lists both these two building blocks, as well as a few intermediate configurations.

Fig. 3.3 shows the power consumption as a function of the total router capacity for various core routers and increasing capacity configurations. As can be seen, Cisco's latest CRS generation (CRS-3) is the most energy efficient. It has been plotted two times, once with 1x100 Gbps port cards installed and once with 14x10 Gbps port cards installed. The latter is more

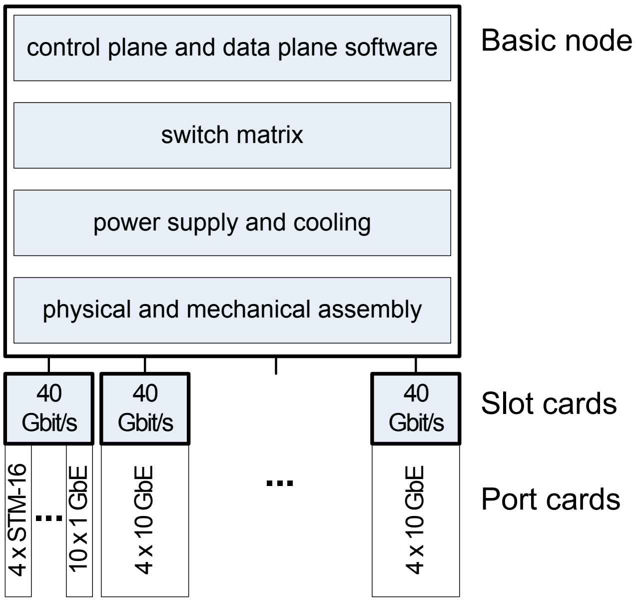

Figure 3.1: *Simplified IP/MPLS router block model (from [6])*

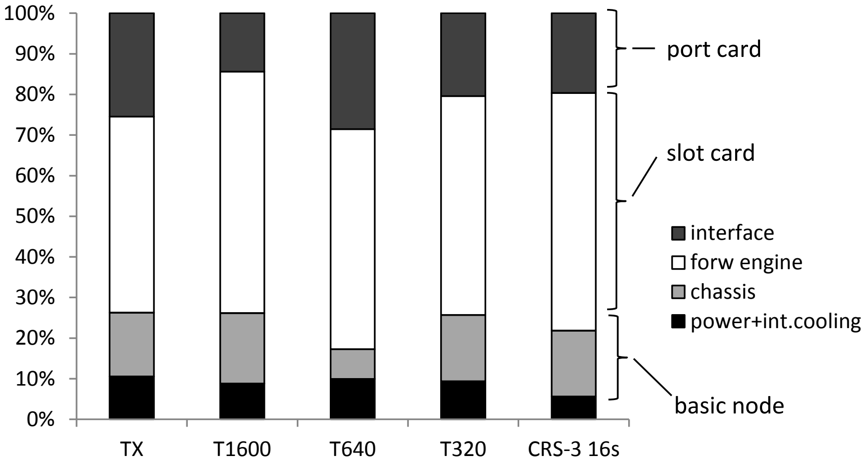

Figure 3.2: *Core router power distribution among the different components*

Table 3.1: *IP/MPLS components*

Basic Nodes: Capacity	*Number of provided slots (slot capacity = 140 Gbps)*	*Power [Watt]*
Line card shelf (2 240 Gbps)	16 slots	2 401
Fabric card shelf (connects max 9 line card shelves)	-	8 100
2 240 Gbps (1 LCS + 0 FCS)	16 slots	2 401
4 480 Gbps (2 LCSs + 1 FCS)	32 slots	12 902
6 720 Gbps (3 LCSs + 1 FCS)	48 slots	15 304
...		
20 160 Gbps (9 LCSs + 1 FCS)	144 slots	29 711
22 400 Gbps (10 LCSs + 2 FCSs)	160 slots	40 212
...		
161 128 Gbps (72 LCSs + 8 FCSs)	1 152 slots	237 686

Slot Cards: Capacity	*Number of provided slots*	*Power [Watt]*
40 Gbps	1 slot/slot	315
140 Gbps	1 slot/slot	401

Port Cards: Port count × Interface Type	*Number of occupied slots*	*Power [Watt]*
16 x PoS STM-16, 80 km	1 slot	122
4 x PoS STM-64, 80 km	1 slot	124
1 x PoS STM-256, 2 km	1 slot	59
8 x 10 Gigabit Ethernet, 40 km	1 slot	79
14 x 10 Gigabit Ethernet, 80 km	1 slot	135
20 x 10 Gigabit Ethernet, 80 km	1 slot	135
1 x 100 Gigabit Ethernet, 10 km	1 slot	135

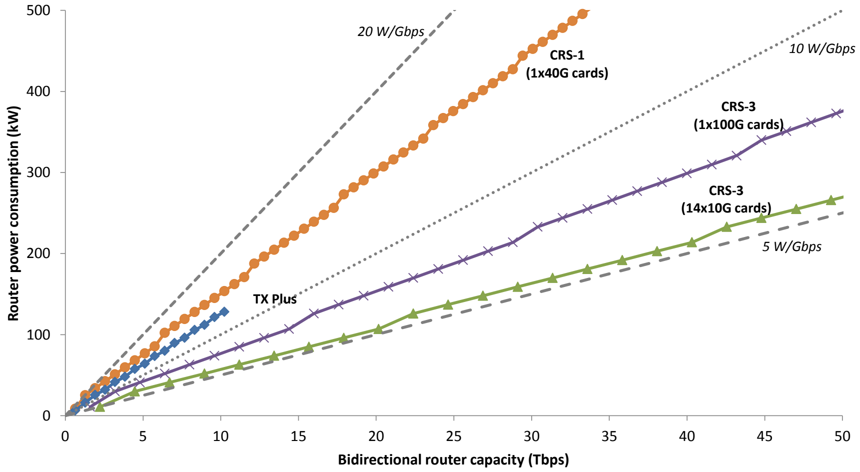

Figure 3.3: *Core router power consumption as a function of the total node capacity, for maximally equipped configurations (full CRS range not shown; all CRS configurations are based on 16-slot CRSs).*

energy efficient because the maximum slot capacity (140 Gbps) is completely used for the same energy consumption. Note that Fig. 3.3 does not show the complete range of the CRS capacity which scales up to 46 Tbps (CRS-1) and 161 Tbps (CRS-3) full duplex.

Based on the values shown in Fig. 3.3, we additionally propose a simplified IP/MPLS layer power value that expresses the power P_{IP} of the node based on the total node capacity C_{IP}:

$$\frac{P_{IP}}{C_{IP}} = 10\ \mathrm{W/Gbps} \tag{3.1}$$

This value is higher than the current achievable CRS-3 energy-efficiency (5.5–7.5 W/Gbps), but seems more reasonable as it implicitly covers suboptimally filled configurations. It is important to note that this value expresses a power efficiency per *equipment capacity*. The actual value might be, and will be, higher (i.e., worse) for real life *throughputs* where the average throughput will be lower than the capacity.

Fixed power-per-port values can be derived from the power-per-node-capacity value given above. For example, a 10G port would consume 100 W.

3.3.2 Ethernet Layer

The Ethernet power consumption is based on two systems: the Cisco Nexus 7018 and the Juniper EX8216. The power consumption values are based on

Table 3.2: *Ethernet layer (bidirectional)*

Type	Power consumption	Power efficiency
Ethernet 1 Gbps port	7 W	7 W/Gbps
Ethernet 10 Gbps port	38 W	3.8 W/Gbps
Ethernet 40 Gbps port	(105 W)	(2.6 W/Gbps)
Ethernet 100 Gbps port	(205 W)	(2.1 W/Gbps)
Ethernet 400 Gbps port	(560 W)	(1.4 W/Gbps)
Ethernet 1 Tbps port	(1100 W)	(1.1 W/Gbps)

Table 3.3: *OTN layer (bidirectional)*

Type	Power consumption	Power Efficiency
OTN 1 Gbps port	7 W	7 W/Gbps
OTN 2.5 Gbps port	15 W	6 W/Gbps
OTN 10 Gbps port	34 W	3.4 W/Gbps
OTN 40 Gbps port	160 W	4 W/Gbps
OTN 100 Gbps port	360 W	3.6 W/Gbps
OTN 400 Gbps port	(1236 W)	(3.09 W/Gbps)
OTN 1 Tbps port	(2794 W)	(2.79 W/Gbps)

the typical power consumption of a maximum configured system, including the power overhead of the chassis and any required control and switch fabric cards.

The values are given in Table 3.2. Power values between brackets represent a projection to higher capacities based on the exponential function for 1 Gbps and 10 Gbps ports.

3.3.3 OTN layer

The Optical Transport Networking (OTN) power consumption is based on confidential information and are approximations. The power consumption values are based on the typical power consumption of a maximum configured system, including the power overhead of the chassis and any required control and switch fabric cards.

The values are given in Table 3.3. Power values between brackets represent a projection to higher capacities based on the exponential function for 40 Gbps and 100 Gbps ports. It is interesting to observe that the power efficiency becomes worse at 40 Gbps. This is probably due to heavy digital signal processing, which is not present in the lower-capacity cards.

3.3.4 WDM layer

WDM component terminology and their associated functions can differ considerably between different vendor and academic documents. To avoid misunderstanding, we first give an overview of the main terminology of the WDM components in this paper. For a more detailed explanation, see [6] or [16].

Transceivers provide full-duplex conversion from/to an electrical signal to/from an optical signal. They are typically commercially available in standardized enclosures such as SFP (1G) and XFP (10G), XENPAK (10G), CFP (100G)[1]. The power consumption of transceivers is usually provided by the power budget of the port card. Therefore, we do not consider individual power consumption of transceivers. *Transponders* are devices that provide bidirectional conversion from one optical wavelength to another, typically from/to a grey (1300 nm) optical signal to a DWDM-band (1500 nm) specific wavelength optical signal. Transponders can be considered as two back-to-back transceivers. The (grey) client side interface typically has limited reach (e.g. up to 2 km, 40 km, or 80 km), whereas the line side interface typically has longer reach (e.g. 200 km, 500 km or 2000 km) given the appropriate amplification (see further). *Muxponders* are similar devices and come typically in an electrical-optical and optical-optical variant. They perform full-duplex time-division multiplexing of lower rate tributary signals into higher rate WDM signals. We treat transponders and optical-to-optical muxponders as one component, since their power consumption (and functionality) is similar for same-rate equipment. *Regenerators* provide 3R (re-timing, re-shaping, re-transmitting) regeneration of optical signals. The distance the signal can travel (span) before regeneration is required depends on the transponder type, data rate, modulation used, fiber quality, etc. A regenerator can be considered as two back-to-back transponders, and is in practice often implemented as such.

Optical Line Amplifiers (OLAs) cater for signal attenuation and are required at a typical interval of 80 km. An OLA system includes an optical amplifier (Erbium Doped Fiber Amplifier (EDFA) or Raman) per fiber and some additional electronics. OLAs are typically unidirectional, however, as all the values in this report are for bidirectional solutions, we give power consumption for bidirectional OLAs that in practice will be composed of two unidirectional OLAs. *WDM terminal systems*, also called WDM (transmission) systems, (de)multiplexes the individual channels (from) into the fiber pair. They consist of a mux/demux, a booster amplifier (to amplify the outgoing optical signal) and a pre-amplifier (to amplify the incoming

---

[1] SFP: Small Form-factor Pluggable, XFP: 10 Gigabit Small Form-factor Pluggable, XENPAK, CFP: C Form-factor Pluggable

optical signal). The WDM terminal is mainly characterized by the number of supported WDM channels (e.g., 40, 80, 96).

Optical switches perform switching of wavelength channels without the need for Optical-Electrical-Optical (OEO) conversion. *Optical Add/-Drop Multiplexers (OADMs)* provide two bidirectional transit fiber ports and are capable of adding-dropping individual wavelengths to a local port. OADMs are characterized by (a) the pass-through capacity at 40 or 80 channels, (b) the percentage of channels that can be added, and (c) the reconfigurability (ROADMs). *OXCs* (optical cross connects) provide more than two bidirectional fiber ports and are capable of cross-connecting wavelength channels. In line with the terminology used in [6], the number of network-side bidirectional fiber ports of an OXC is known as the degree. This does not include the add/drop fiber ports which we label as the add/drop degree. For degree-2 nodes, ROADMs can be used, for multi-degree switching OXCs are used (which can be implemented in practice by combining a number of ROADMs). Different technologies can be used for implementing optical switches, e.g., Microelectromechanical Systems (MEMS) or liquid crystal-based wavelength selective switches. Unfortunately, the underlying technology was unclear for the provided values. It is probably MEMS though.

Dynamic Gain Equalizers (DGEs) and Dispersion-compensating Fibers (DCFs), which provide signal conditioning, are not considered. They are either passive devices with negligible indirect power consumption impact on the other components, or consume negligible power.

The values listed in Table 3.4 are the proposed values for the various WDM components. These values are based on a generalization of data sheet power consumption values of a wide number of components [5]. The values between brackets indicate projected values. Node degree d is the number of network-side bidirectional fiber ports. The add/drop degree a is the number of add/drop bidirectional fiber ports, potentially ranging from 0 to d. Note that the transponder values provided in Table 3.4 are for non-coherent transponders. Values for coherent transponders will be higher, but no public values are available yet. Coherent transponders are used to increase the transmission distance at higher bandwidths.

3.4 Analytical power consumption model

In this section we propose a simplified analytical power consumption model for the various layers. The model is given first (Section 3.4.1). The details on how the model is constructed follow (Section 3.4.2).

Table 3.4: *WDM components (bidirectional)*

Type	Remarks	Power [Watt]
Transponder/Muxponder 2.5G	Per channel pair, includes overhead. All non-coherent transponders.	25 W
Transponder/Muxponder 10G		50 W
Transponder/Muxponder 40G		100 W
Transponder/Muxponder 100G		(150 W)
Transponder/Muxponder 400G		(300 W)
Transponder/Muxponder 1T		(500 W)
Regenerator xG	Per channel pair, includes overhead	2 × transp. xG
OLA, short span 2 km	Per fiber pair (!), includes overhead	65 W
OLA, medium span 40 km		65 W
OLA, long span 80 km		110 W
OLA, very long span 120 km		120 W
WDM terminal, 40 channels	Per fiber pair, includes mux/demux, pre- and booster amplifier, and overhead	230 W
WDM terminal, 80 channels		240 W
ROADM, 40 channels, 100%	Per node, includes mux/demux, pre- and booster amplifier, and overhead	450 W
ROADM, 80 channels, 50%		550 W
ROADM, 80 channels, 100%		600 W
OXC, 40 channels, node degree d, add/drop degree a	Per node, includes mux/demux (for add/drop), pre- and booster amplifier, and overhead	$d\times$ 85 W + $a\times$ 50 W + 150 W
OXC, 80 channels, node degree d, add/drop degree a		$d\times$ 85 W + $a\times$ 100 W + 150 W

3.4.1 Model

The total power P_{core} [Watt] in an optical multilayer core network is the sum of the power consumption in the constituting layers:

$$P_{core} = P_{ip} + P_{ethernet} + P_{otn} + P_{wdm} \tag{3.2}$$

with

$$P_{wdm} = P_{optsw} + P_{transponders} + P_{amplifiers} + P_{regeneration} \tag{3.3}$$

The power consumption for each layer can be written as a function of the average IP demand $\overline{D_C}$, a power efficiency P/C value for that layer, and the hop count H for each layer:

$$\begin{aligned}
P_{ip} &= \eta_c \cdot \eta_{pr} \cdot N_d \cdot \overline{D_C} \cdot \left[\frac{P_{IP}}{C_{IP}} \cdot 2 \cdot \left(\frac{1}{\eta_{pr}} + H \right) \right] \\
P_{ethernet} &= \eta_c \cdot \eta_{pr} \cdot N_d \cdot \overline{D_C} \cdot \left[\frac{P_{ETH}}{C_{ETH}} \cdot 2 \cdot \left(\frac{1}{\eta_{pr}} + H \right) \right] \\
P_{otn} &= \eta_c \cdot \eta_{pr} \cdot N_d \cdot \overline{D_C} \cdot \left[\frac{P_{OTN}}{C_{OTN}} \cdot 2 \cdot \left(\frac{1}{\eta_{pr}} + H \right) \right] \\
P_{optsw} &= \eta_c \cdot \eta_{pr} \cdot N_d \cdot \overline{D_C} \cdot \left[\frac{P_{OXC}}{C_{OXC}} \cdot 2 \cdot H \right] \\
P_{transponders} &= \eta_c \cdot \eta_{pr} \cdot N_d \cdot \overline{D_C} \cdot \left[\frac{P_{TR}}{C_{TR}} \cdot 2 \cdot H \right] \\
P_{amplifiers} &= \eta_c \cdot \eta_{pr} \cdot N_d \cdot \overline{D_C} \cdot \left[\frac{1}{f} \frac{P_{OLA}}{C_{OLA}} \cdot \left\lfloor \frac{\alpha}{L_{amp}} \right\rfloor \cdot H \right] \\
P_{regeneration} &= \eta_c \cdot \eta_{pr} \cdot N_d \cdot \overline{D_C} \cdot \left[\frac{P_{RE}}{C_{RE}} \cdot \left\lfloor \frac{\alpha}{L_{regen}} \right\rfloor \cdot H \right]
\end{aligned} \tag{3.4}$$

The symbols with description and reference values are listed in Table 3.5 Remarks:

- The *power efficiency values* P/C have been determined by dividing the power values from Section 3.3 by the capacity of the corresponding component. Exemplary values are given for 2.5G, 10G and 100G equipment.
- The *booster and pre-amplifier* power consumption is accounted for optical switching instead of the amplifiers, see further.
- The factor η_c accounts for *cooling* and *facilities overhead* power consumption in telecom centers. This overhead is commonly characterized by the Power Usage Effectiveness (PUE) [17]. The PUE is the

Table 3.5: *Symbols and values*

Quantity	Symbol	Value (2.5G)	Value (10G)	Value (100G)
Efficiency, IP/MPLS core router	P_{IP}/C_{IP}	10 W/Gbps	10 W/Gbps	10 W/Gbps
Efficiency, Ethernet	P_{ETH}/C_{ETH}	1.3 W/Gbps	3.8 W/Gbps	2.1 W/Gbps
Efficiency, OTN	P_{OTN}/C_{OTN}	6.0 W/Gbps	3.4 W/Gbps	3.6 W/Gbps
Efficiency, Optical switching,	P_{OXC}/C_{OXC}			
ROADM, 100%, 40 ch.		2.25 W/Gbps	0.56 W/Gbps	0.06 W/Gbps
OXC, degree=3, 40 ch.		1.85 W/Gbps	0.46 W/Gbps	0.05 W/Gbps
Efficiency, transponder	P_{TR}/C_{TR}	10 W/Gbps	5 W/Gbps	1.5 W/Gbps
Efficiency, optical line amplifier, long span (40 ch)	P_{OLA}/C_{OLA}	1.1 W/Gbps	0.27 W/Gbps	0.03 W/Gbps
Efficiency, regenerators	P_{RE}/C_{RE}	20 W/Gbps	10 W/Gbps	3 W/Gbps
Provisioning factor for protection	η_{pr}	2		
Provisioning factor for cooling and facilities overhead (=PUE)	η_c	2		
Average layer hop count	H	Depends on network topology, traffic demands and routing		
Total number of IP/MPLS demands	N_d	Given by the traffic matrix		
Average demand capacity	$\overline{D_C}$	Given by the traffic matrix		
Average fiber filling (% of used channels in fiber)	f	Depends on network topology, traffic demands, established lightpaths and routing		
Average (lightpath) link length	α	Given by the network topology		
Optical amplification span length	L_{amp}	80 km		
Optical regeneration length	L_{regen}	1500 km		

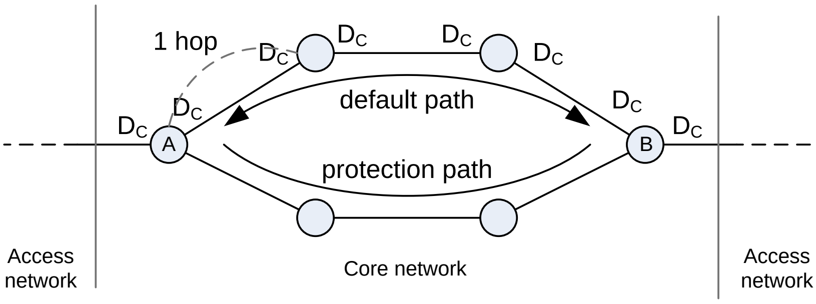

Figure 3.4: *Required router ports for one 1+1 protected demand*

ratio of the total amount of power consumed over the useful power consumed, and typically has a value of 2 [18]. In highly optimized and efficiently cooled data centers, lower PUE values are possible, but this is not yet commonplace. The subscript c has been chosen to be in line with the terminology used in [12].

- The factor η_{pr} accounts for *traffic protection*, and equals 2 for 1+1 protection. For unprotected traffic the value would be 1.
- The average IP/MPLS-layer *hop count* H is the number of hops in the respective layer averaged over all traffic demands. For a given topology, the hop count will depend on such aspects as the routing algorithm, link weights, etc. For the equations to be valid, each hop in the IP/MPLS layer means the termination of a lightpath.

3.4.2 Explanation

Power consumption in the *IP/MPLS* layer is calculated according to the number of router ports required for supporting a single bidirectional (i.e., full-duplex) demand with capacity D_C between nodes A and B, see Fig. 3.4.

So, the resulting IP/MPLS capacity T_C (in [Gbps]) required for this single demand is given by:

$$T_C = D_C + \eta_{pr} \left(2 \cdot H \cdot D_C\right) + D_C = 2 \cdot D_C \cdot (1 + \eta_{pr} \cdot H) \qquad (3.5)$$

As we can see, it is a function only of the demand capacity D_C, the number of routing hops H and the protection factor $\eta_{pr} = 2$. Note that we assume the number of hops in the protection path to be equal to the number of hops in the default path.

Thus, if we assume an average demand capacity $\overline{D_C}$, the required total IP/MPLS capacity T_{IP} (in [Gbps]) is given by multiplying with the total number of demands N_d:

$$T_{IP} = N_d \cdot \overline{T_C} \tag{3.6}$$

The power consumption in the IP/MPLS layer P_{IP} is the total capacity T_{IP} multiplied by the bidirectional (or full-duplex) power efficiency E_{IP} of this layer.

The power efficiency E_{IP} of the IP/MPLS layer is determined by the power consumption of the router (i.e., basic node equipped with slot and port cards) for a given capacity (P_{IP}/C_{IP}) and any additional external overhead power, indicated by the factor η_c. P_{IP}/C_{IP} = 10 W/Gbps is the value proposed in Section 3.3.1. The overhead factor η_c will typically be 2 or less (for newer premises).

Thus, we get for the power consumption in the routing layer (in [Watt]):

$$P_{routing} = E_{IP} \cdot T_{IP} = \left(\eta_c \cdot \frac{P_{IP}}{C_{IP}}\right) \cdot \left(N_d \cdot \overline{D_C} \cdot 2 \cdot (1 + \eta_{pr} \cdot H)\right) \tag{3.7}$$

For the *Ethernet*, and *OTN* we deduce identically.

For the *transponders* and the *optical switching devices* we deduce identically, with the exception that we do not account for a long haul transponder at the access network sides.

For the *OLAs* we have (see Fig. 3.5):

$$T_C = \eta_{pr} \cdot \left\lfloor \frac{\alpha}{L_{amp}} \right\rfloor \cdot H \cdot D_C \tag{3.8}$$

A fiber filling factor f is added in the final equation of Eq. (3.4) to account for suboptimal usage of fiber channels. Note that we did not account for the booster and pre-amplifiers in Eq. (3.8), because we consider them to be part of the optical switching devices. However, if required they could be accounted for by slightly modifying Eq. (3.8) to:

$$T_C = \eta_{pr} \cdot \left\lfloor 2 + \frac{\alpha}{L_{amp}} \right\rfloor \cdot H \cdot D_C \tag{3.9}$$

For the *regeneration*, the idea is identical to the OLAs. The number of regenerators per demand is approximated by the factor $\left\lfloor \frac{\alpha}{L_{regen}} \right\rfloor \cdot H$. However, if the link lengths α are in the same order of the regeneration length L_{regen} (taken to be 1500 km), the approximation will be rather crude. An alternative approach would be to replace the earlier factor with a more general regeneration factor η_r expressing the number of regenerations per demand, which could be estimated by a more accurate heuristic.

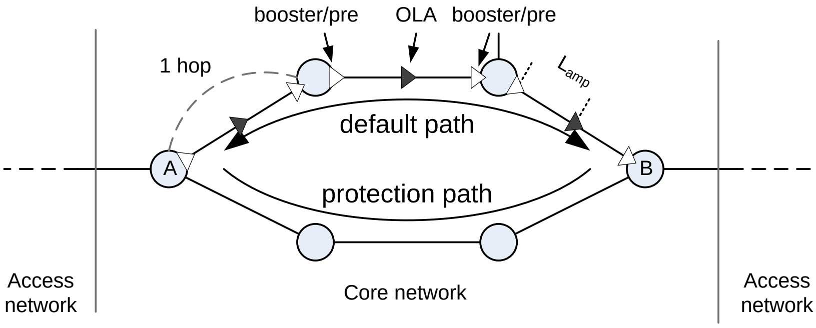

Figure 3.5: *Required optical line amplifiers for one 1+1 protected demand*

3.4.3 Comparing with earlier analytical models

It is useful to compare our model with the models in [11] and [12]. If we look at the power efficiency equation for the IP routing layer, given by equation (13) in [11] and the Table III Long Haul subnetwork PR/CR term in [12], and in both cases ignoring the factor for future provisioning, these models have:

$$E_{IP} = \eta_c \cdot \eta_{pr} \cdot (H+1) \cdot \left(\frac{P_{IP}}{C_{IP}}\right) \tag{3.10}$$

However, our model has:

$$E_{IP} = \eta_c \cdot 2 \cdot (\eta_{pr} \cdot H + 1) \cdot \left(\frac{P_{IP}}{C_{IP}}\right) \tag{3.11}$$

The apparent difference in factor 2 comes from the fact that the two earlier models consider unidirectional (i.e., half-duplex) demands but use a bidirectional P_{IP}/C_{IP} value; as such the bidirectional value eliminates the factor 2. We feel this is confusing, and thus consider both bidirectional (full-duplex) demands *and* efficiencies. The difference in application of the protection factor is because of a simplification by the existing models where the protection capacity is accounted both on the network side and the client side. For example, with $\eta_{pr} = 2$ (e.g. for a 1+1 protection scheme) the add/drop traffic is counted twice. In practice there will be only one add/drop port at the client side (see e.g. [19]), and is the approach we have taken in our model. So, the models are very similar, with the only difference being the protection scheme more accurately modeled in this work.

3.5 Evaluation and case study

In this section, we show how the power consumption values listed in Section 3.3 and the analytical power model from Section 3.4 can be used to calculate the power consumption of a network. This also allows us to evaluate the analytical power consumption model.

3.5.1 Cases considered

We consider two different networks to which we apply a number of traffic matrices: the pan-European network and the American NSFNET network.

To calculate the power consumption associated with these demands, we use two different calculation methods (via simulation, and via the analytical hop count model), and in addition consider two separate scenarios (a router bypass scenario, and a non-bypass scenario).

In the next subsections, we provide more details on each of these cases.

3.5.1.1 Network topologies

We consider two different test networks (see Fig. 3.6) to calculate and evaluate the power consumption:

- the *pan-European core network* is based on the Géant research network [20], but has been modified to represent a commercial transport network (for example, to protect against single link failures, the topology has been modified so that each node is at least connected to two other nodes). We have used the DICONET pan-EU topology [21], which contains 34 nodes and 54 WDM links.
- *NSFNET*, a US network based on a former NSF network topology which has been used in many studies, e.g. [22]. It consists of 14 nodes and 21 WDM links.

The network parameters are summarized in Table 3.6.

3.5.1.2 Network traffic demands

For our case study and evaluation, we apply various traffic matrices, summarized in Table 3.7.

For the pan-EU network we consider: (a) a gravity traffic matrix where nearby nodes have larger demands, thus closer resembling real life demands [21], (b) a random fully-meshed traffic matrix, and (c) a uniform fully-meshed traffic matrix where all demands are equal.

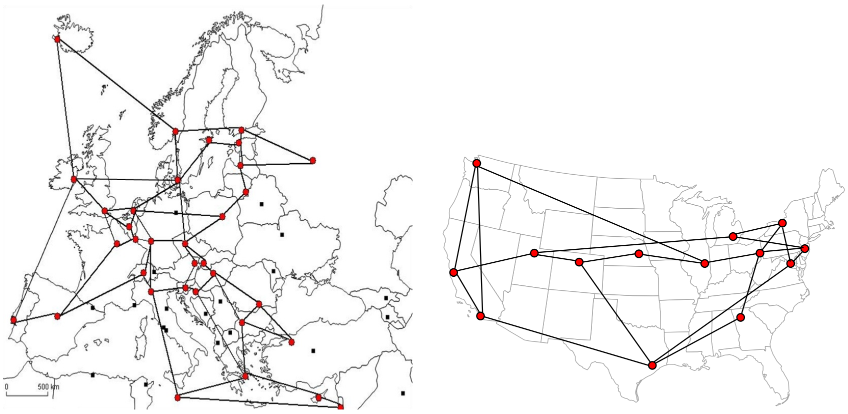

(a) Pan-EU network **(b)** NSFNET network

Figure 3.6: *IP topologies of the test networks*

Table 3.6: *Network topology parameters*

Parameter	Pan-European network	NSFNET
Number of nodes	34	14
Number of links	54	21
Average node degree	3.09	3
Average link length	753 km	1083 km
Minimum link length	67 km	260 km
Maximum link length	2361 km	2840 km

Table 3.7: *Traffic matrices*

Parameter	Pan-EU, gravity	Pan-EU, random	Pan-EU, uniform	NSFNET, random
Number of IP/MPLS demands	367	561	561	91
Actual hop count (by simulation, see Section 3.5.1.4)	4.1	4.6	4.6	2.9
Estimated hop count (see Section 3.5.1.4)	3.83	3.83	3.83	2.45

For NSFNET we only consider a random fully-meshed traffic matrix.

In all the cases above, we scale up the traffic demands, so that we load the network with 10 different traffic matrices ranging from 2.5 to 100 Gbps of average traffic demand.

3.5.1.3 Node architecture and (non-)bypass scenarios

For both networks, we consider the architectural setup shown in Fig. 3.7. Other architectures are possible, e.g. IP-over-OTN-over-WDM see e.g. [6].

In the IP/MPLS layer, a core router is equipped with line cards, providing short reach interfaces. The granularity for the interfaces differs: the access or client-side traffic connects to the router using 1 Gbps interfaces, the core network side channel interfaces are all 10 Gbps interfaces. Note that, depending on the demand capacity, one or more interfaces will be required per demand.

In the WDM layer, long reach transponders provide a Dense Wavelength Division Multiplexing (DWDM) optical signal, which is switched using an OXC to the correct link. A mux/demux aggregates up to 40 channels on a fiber. For each link, we assume an unlimited numbers of fibers to be available. A booster and pre-amplifier amplify all signals in a fiber pair respectively upon leaving or entering a node. An inline amplifier is placed every 80 km. For link lengths longer than the regenerator span, taken to be 1500 km, the signal is switched by the OXC to pass through a regenerator. The regenerator itself is composed of 2 back-to-back transponders.

These architectural assumptions are summarized in Table 3.8.

With this architecture in mind, we consider two different scenarios for calculating the power consumption:

- A *non-bypass scenario*, where all traffic in the node—both the traffic that starts or ends in the node, as well as the transit (bypass) traffic—is processed by the core router. This provides opportunity for the IP router to groom—i.e., bundle traffic demands from different sources

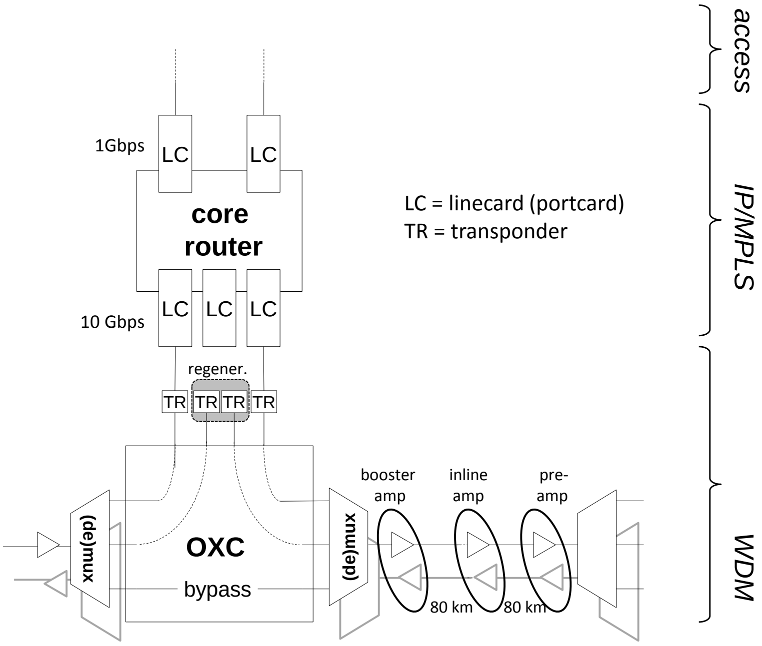

Figure 3.7: *Network node and link architecture*

Table 3.8: *Network architectural parameters*

Parameter	Value
Optical amplification span L_{amp}	80 km
Regenerator span L_{regen}	1500 km
Channels per fiber	40
Channel capacity	10 Gbps
Protection	1+1
Node client-side capacity interface granularity	1 Gbps
Node network-side capacity interface granularity	10 Gbps

destined for the same outgoing link. This assures that optical channels can be optimally filled.

- An *optical bypass scenario*, where a dedicated lightpath (channel) is set up from source node to destination node. By doing so, we create a new, modified IP topology which we call the virtual topology. This way, the transit (bypass) traffic destined for another node does not have to be handled by the IP router, and consequently not have to be converted from the optical to the electronic domain and back to the optical domain. On the other hand, if a source-destination traffic demand is smaller than the available channel capacity, the channel will not be optimally used, resulting in a higher number of channels and equipment required. Note that our optical bypass scenario is the extreme case of applying optical bypass. More intermediate cases would consist of optical multi-hop bypass.

Furthermore, we assume that the network provides 1+1 protection, which means that for each demand two link-disjoint IP connections or lightpaths are set up. If one path fails, the traffic is still available without interruption over the other path.

3.5.1.4 Calculation methods

We use two different methods to calculate the power consumption in the networks, of which we then compare the resulting values. In both cases, we assume a PUE of 2.

Using simulation to dimension the network The first method is based on dimensioning the network via simulation, that is, calculating for each traffic demand the path that will be followed across all nodes, and subsequently determine the equipment required. By multiplying the equipment count with its respective power consumption, the total power is

Table 3.9: *Dimensioning via simulation power values*

Parameter	Value	Unit
IP router efficiency	10	W/Gbps
Transponder (10G, bidirectional)	50	W
Regenerator (10G, per bidirectional channel)	100	W
OLA, long span 80 km	110	W
OXC, average node degree $\bar{d}$, with add/drop degree $a = d$	($\bar{d}\cdot$ 135 + 150)	W
Power usage effectiveness (PUE)	2	-

Table 3.10: *Analytical power model valuess*

<table>
<tr><th>Parameter</th><th>Symbol</th><th>Value (non-bypass)</th><th>Value (bypass)</th></tr>
<tr><td>IP router efficiency</td><td>P_{IP}/C_{IP}</td><td colspan="2">10 W/Gbps</td></tr>
<tr><td>Transponder efficiency (10G)</td><td>P_{TR}/C_{TR}</td><td colspan="2">5 W/Gbps</td></tr>
<tr><td>Regenerator efficiency (10G)</td><td>P_{RE}/C_{RE}</td><td colspan="2">10 W/Gbps</td></tr>
<tr><td>OLA efficiency</td><td>P_{OLA}/C_{OLA}</td><td colspan="2">0.27 W/Gbps</td></tr>
<tr><td>OXC efficiency (40 10G-channels)</td><td>P_{OXC}/C_{OXC}</td><td colspan="2">0.46 W/Gbps</td></tr>
<tr><td>Average IP/MPLS hop count</td><td>H</td><td>see text</td><td>1</td></tr>
<tr><td>Average hop count optical sw.</td><td>H'</td><td colspan="2">see text</td></tr>
<tr><td>Provision. factor for protection</td><td>η_{pr}</td><td colspan="2">2</td></tr>
<tr><td>Provision. factor for cooling (PUE)</td><td>η_c</td><td colspan="2">2</td></tr>
<tr><td>Number of IP/MPLS demands</td><td>N_d</td><td colspan="2">see text</td></tr>
<tr><td>Average demand capacity</td><td>$\overline{D_C}$</td><td colspan="2">see text</td></tr>
<tr><td>Average fiber filling</td><td>f</td><td colspan="2">100%</td></tr>
<tr><td>Average (lightpath) link length</td><td>α</td><td>see text</td><td>see text</td></tr>
</table>

determined. We route the demands using a shortest cycle algorithm (to provide 1+1 protection) and wavelengths are selected following a first-fit wavelength assignment algorithm [23]. The power values used are summarized in Table 3.9. For all components we use the power values listed in Section 3.3. Because of simulation tool constraints we generalized on the OXC power consumption and calculate an average OXC power consumption value based on the average node degree of the network (see Table 3.6).

Using the analytical power model The second method uses the *analytical power model* proposed in Section 3.4. This is less accurate than the simulation approach, but has the advantage of being trivial to compute, as it only requires filling in the parameters in the equations. The values used are listed in Table 3.10.

The parameter values were determined as follows:

- The IP core *router*, *transponder*, *regenerator* and *amplifier* efficiency values are as earlier defined.

- The *OXC* efficiency P_{OXC}/C_{OXC} for one demand is approximated by the OXC power consumption for the average node degree $\bar{d}$, divided by the total capacity of the OXC. Thus, for our 40-channel OXC, we get:

  $$\frac{P_{OXC}}{C_{OXC}} = \frac{\left(150[\mathrm{W}] + \bar{d} \cdot (85[\mathrm{W}] + 50[\mathrm{W}])\right)}{40 \cdot 10[\mathrm{Gbps}] \cdot \bar{d}} \tag{3.12}$$

  For the pan-European and NSFNET network, the average node degree $\bar{d}$ is 3.09 and 3 respectively (see Table 3.6). Thus, the value for both networks is almost identical, and approximates to 0.46 W/Gbps, in line with Table 3.5.

- Following the network global expectation model proposed in [7], the *hop count* H in a uniform network can be approximated by the following equation, with N the total number of nodes in the network and L the number of bidirectional links in the network:

  $$H = \sqrt{\frac{N-2}{\frac{2L}{N}-1}} \tag{3.13}$$

  For the *non-bypass scenario*, for the pan-European network we have N=34 and L=54, which gives H=3.83, whereas for the considered traffic demands routed by the shortest cycle algorithm as described above, the hop count is 4.1 and 4.6 (see Table 3.7). As our analytical power model scales linearly with the hop count, the error on the result will be equally large. As such, to evaluate the proposed power model fairly, we will use the actual hop count as determined by dimensioning the network via simulation with a given traffic matrix. These values are listed in Table 3.7, both for the pan-EU network and NSFNET.

  For the *bypass scenario* the hop count H is 1, as we have created a new virtual IP topology where direct source-destination lightpaths are set up. However, the *hop count for the optical switching* H', remains identical to the non-bypass scenario hop count, as each connection traverses an OXC regardless of the scenario.

- The *number of demands* is directly available from the traffic matrix, as well as the average demand capacity.

- The *average fiber filling* is estimated to be 100%, which will be a good approximation for large demands.

- The *average (lightpath) link length* is given directly by the network topology (see Table 3.6). Again, it could also be estimated; the network global expectation model [7] provides an approximation based on the geographic area A covered by the network $\alpha = \sqrt{A}/\left(\sqrt{N}-1\right)$. For an estimated area of the pan-European network of 3000×3300 km^2, this would give α = 653 km, which gives only a 13% difference from the actual value of 753 km.

  As for the bypass scenario we have a hop count equal to one, the lightpath link length equals the sum of lengths of all the fibers that the lightpath is traversing.

3.5.2 Results

3.5.2.1 Model evaluation

Fig. 3.8 shows the result of applying the various traffic matrices (see Section 3.5.1.2) to the pan-EU and NSFNET networks. The charts map the power consumption with average traffic demand increasing to up to ten times the channel and port capacity (10 Gbps). The solid lines represent the power consumptions as calculated by the simulation approach. The dashed lines indicate the result from the analytical power model. The upper lines are the power consumption for the non-bypass scenario, while the lower lines are for the optical bypass scenario.

We make the following observations:

- The analytical power model approaches very well the simulation result (Fig. 3.8). In the non-bypass scenario, for high demands (relative to the channel capacity) the approximation converges to 97% for the pan-EU network and 93% for NSFNET. Note that, as explained in Section 3.5.1.4, part of this good approximation is because we used the actual hop count value in our analytical model, as determined through simulation, instead of a heuristic to approximate it.

- The estimation is very good for all layers except the regeneration (Fig. 3.9). This is the result of the crude approximation made for the number of regenerations per demand (see Section 3.4.2). For the non-bypass scenario the mathematical flooring of the average link length over the regeneration length gives zero, resulting in zero power for the regeneration. On the other hand, for the bypass scenario, the regeneration estimate is too high.

- The crude regeneration estimation is also the reason for the NSFNET approximation to be lower than the pan-EU approximation. As the total power consumption for NSFNET is much lower (because of the lower number of nodes, and thus demands, for an equal average traffic demand), and because of the longer link lengths (see Table 3.6), the influence of the regeneration estimation error is relatively larger, see Fig. 3.9(c) and (d).
- For the optical bypass scenario the approximation is good for high demands. However, it does fall short for low demands, as clearly shown in Fig. 3.10. This is no surprise, as for traffic demands below the network interface capacity (i.e., below 10 Gpbs) the model does not take into account the suboptimal used interfaces, thereby overestimating the router efficiency; the underestimation is much worse than for the non-bypass scenario because in the latter the grooming dampens the sub-optimality.

3.5.2.2 Component power consumption distribution

If we look more in detail to the distribution over the different components (Fig. 3.11), we see that the largest share of power consumption is concentrated in the IP router. The transponders are the second major contributor. This follows also directly from the difference in efficiency (for 10G equipment, we defined P_R/C_R to be the double of P_{TR}/C_{TR}, see Table 3.5). This is also in line with earlier findings such as in [3], however, the figures differ slightly. For example, [3] attributes 90% to the routers and 5% to the transponders. This is due to the very high power consumption (1000 W) assumed for an IP router port.

Amplification and regeneration power consumption only becomes relevant in the bypass scenario. For the amplifiers this is only because of the reduction of the total power consumption, as the absolute amplifier power consumption remains constant; for the regenerations this relative increase is in addition caused by the longer link lengths, see Section 3.5.1.4.

Fig. 3.12 shows that indeed the IP router and transponder power consumption has decreased for the bypass scenario, and that the amplifier and OXC power consumption remains the same in both scenarios. A lot of skipped router hops were replaced by regenerator hops, which is shown in the increased regeneration power consumption.

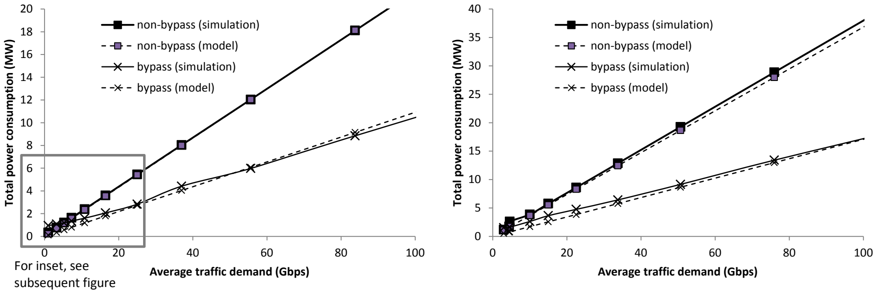

(a) pan-EU (gravity demands)

(b) Pan-EU (random demands)

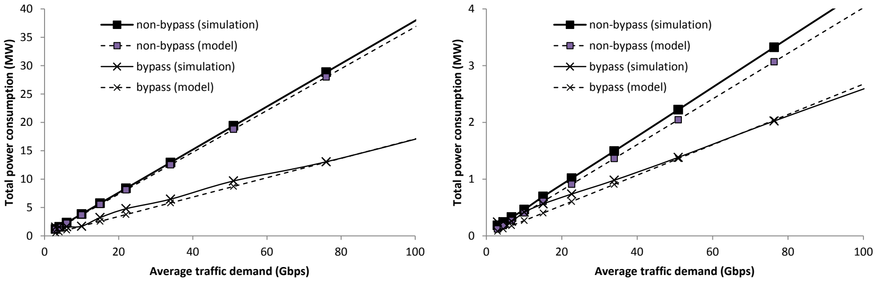

(c) Pan-EU (uniform demands)

(d) NSFNET (random demands)

Figure 3.8: *Power consumption with increasing traffic demands (gravity, random and uniform) for the pan-EU and NSFNET network.*

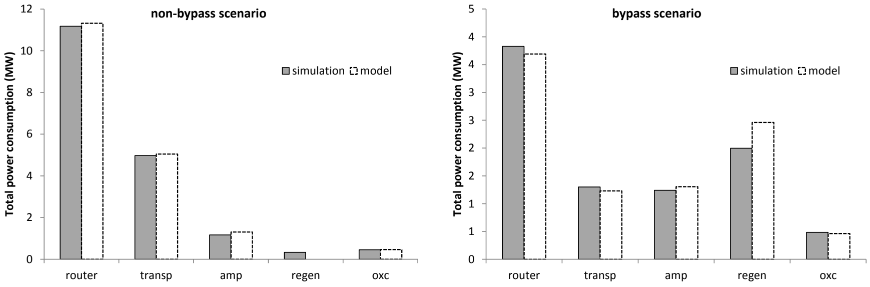

(a) Pan-EU, gravity matrix, non-bypass **(b)** Pan-EU, gravity matrix, bypass

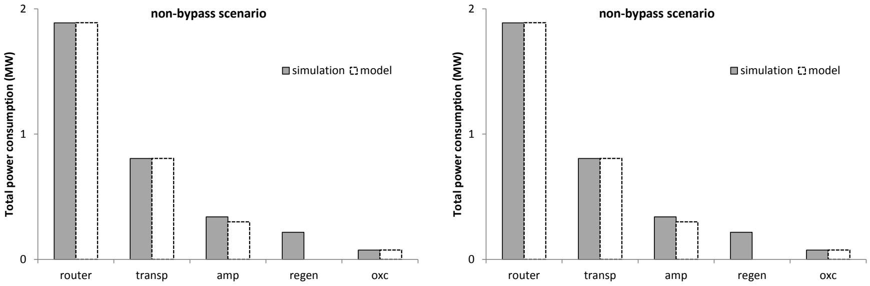

(c) NSFNET, random matrix, non-bypass **(d)** NSFNET, random matrix, bypass

Figure 3.9: *Comparison of the component power consumptions as calculated by the simulation approach (grey bars), and the power model approach (white bars). Average traffic demand equals 80 Gbps for all cases.*

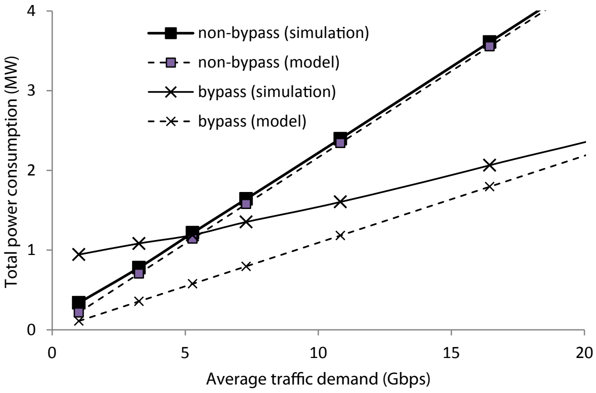

Figure 3.10: *Detail of inset in Fig. 3.8*

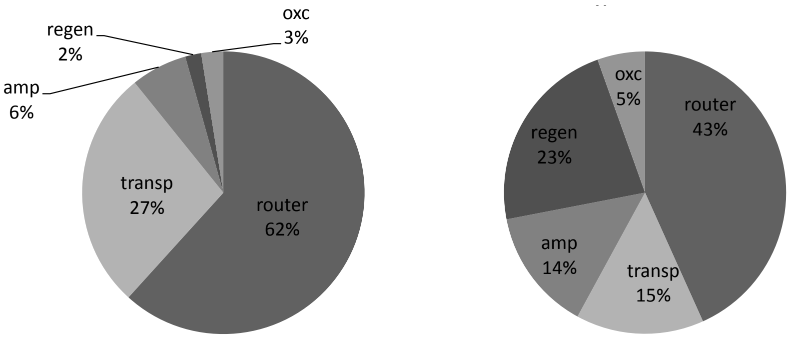

(a) Non-bypass scenario **(b)** Bypass scenario

Figure 3.11: *Relative component power consumption for the pan-EU network at 80 Gbps average demand (simulation results)*

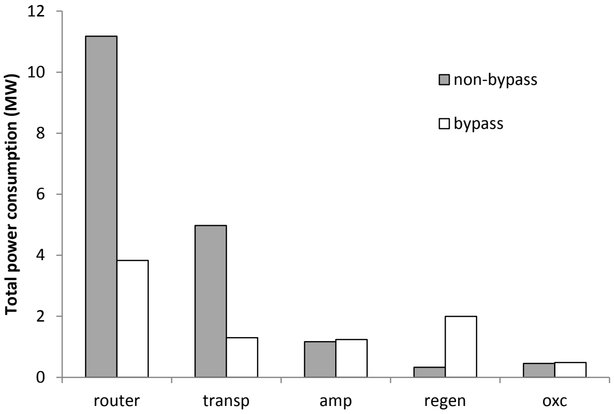

Figure 3.12: *Component power consumption for the pan-EU network (gravity matrix, 80 Gbps)*

The OXC power consumption is negligible in both scenarios.

While Fig. 3.11 and Fig. 3.12 show only the case for the pan-EU network with gravity demands, the results for the 3 other cases are very similar. The longer link lengths in the NSFNET slightly increase the relative contribution of the amplifier and regeneration power consumption .

3.5.2.3 Savings from optical bypass

As already shown in earlier figures, the optical bypass scenario clearly provides potential for significant power savings over the non-bypass scenario,

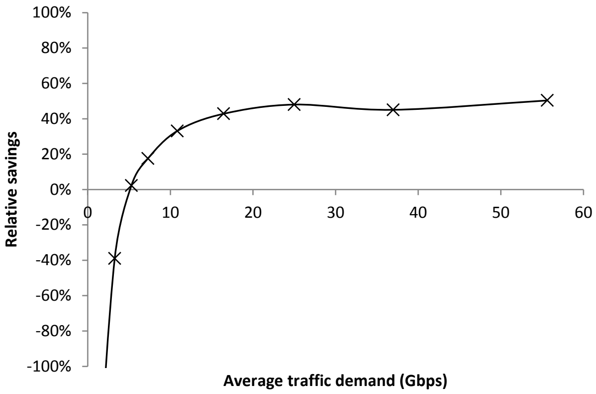

Figure 3.13: *Relative savings of bypass over non-bypass (pan-EU network, gravity matrix)*

but not under all circumstances. In Fig. 3.13 the relative savings of the bypass scenario over the non-bypass scenario for the pan-EU network are mapped.

For low demands savings are negative, i.e., optical bypass consumes more energy. This is because for optical bypass, at least one dedicated optical channel is required for each source-destination demand. As at the network side we only have 10 Gbps interfaces available, for demands below 10 Gbps the channels are not optimally filled. This is less the case for the non-bypass scenario where all the traffic is 'pulled' up into the IP routing layer: the router can groom all traffic demands for the same outgoing links, thereby optimally filling the channels, saving on the number of 10 Gbps interfaces and subsequently power consumption.

With rising traffic demands (from around 4 Gbps of average traffic demand), the bypass strategy starts to pay off, consuming less energy. The power consumption of the bypass strategy initially rises slower than for the non-bypass strategy (see also Fig. 3.10). This is because the underutilized 10 Gbps channels can carry the additional traffic demands at almost no energy increase, whereas for the non-bypass this is not the case.

For high demands—i.e., higher than the channel and interface capacity, which is 10 Gbps—savings converge to about 50%. The slight drop around 37 Gbps is because of the coincidental large number of 11 Gbps demands in the traffic matrix, which in the bypass scenario results in one of the two required channels being suboptimally filled. It is important to point out that the 50% value is no magic number. As shown in [3], the maximum energy savings achieved depend on the size of the network (in terms of nodes). For a network with similar connectivity, gains will be lower for smaller networks, and higher for larger networks. This is because for larger networks the chance of establishing longer lightpaths increases, bypassing more intermediate nodes, and thus saving on router interfaces. This is confirmed by our findings, which indicate that for the NSFNET network (only 14 nodes, with the pan-EU network having 34 nodes) the savings converge to around 40%.

3.6 Conclusions

This paper has two main objectives: (a) provide traceable and well-defined power consumption estimates for optical multilayer network equipment, and (b) provide an analytical power consumption model that avoids the need for network dimensioning, for example via simulation.

The equipment power consumption values are defined for reference in Table 3.1, Table 3.2, Table 3.3 and Table 3.4. They represent typical values (as opposed to maximum power consumption values), include chassis and controller overhead, and are for bidirectional (full-duplex) equipment and traffic. We note that our values for optical amplifiers are typically higher than values used in earlier academic works. In contrast, our IP router power consumption values are typically lower, partly due to technical advances in power efficiency. All values are best-effort representations for the current situation. Suggestions for extrapolations to future values and efficiencies are mentioned in e.g. [12].

The analytical power model we propose is mainly based on the average hop count and aligns nicely with earlier work such as [11] and [12]. It provides a good approximation to the power consumption obtained by simulation, *if* the hop count is correctly determined or estimated, and if the equipment capacity (interfaces, channels, ...) is not over-provisioned for the actual demands (e.g. when employing optical bypass). As such, research into more accurate hop count estimation for a given network, traffic demand pattern and routing policy would be useful. The estimation for the regeneration power consumption is less accurate, and would also benefit from a more accurate heuristic to estimate the amount of regenerations per demand.

Our analysis confirms that for current networks the main share of the power consumption—in the order of 60%—is in the IP/MPLS layer, although we found the share to be less than in earlier publications. Transponders are second in power consumption, in the order of a fifth or a quarter of the total power consumption. OXC power consumption is currently negligible.

Optical bypass is a valuable technique to save power, in our exemplary network up to 50%. Savings however depend on the size of the network, and require optimally used interfaces.

Acknowledgments

The authors would like to thank the STRONGEST and TREND members for their valuable feedback and input. The work described in this paper was carried out with the support of the BONE project (Building the Future Optical Network in Europe) and the TREND project (Towards Real Energy-efficient Network Design), both a Network of Excellence funded by the European Community's Seventh Framework; the IBBT-project GreenICT and the STRONGEST project funded by the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement n 247674.

References

[1] M. Gupta and S. Singh. *Greening of the Internet*. In Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM), pages 19–26, Karlsruhe, Germany, 2003. ACM. doi:10.1145/863955.863959.

[2] M. Webb et al. *SMART2020: Enabling the low carbon economy in the information age*. The Climate Group, 2008. Available at http://www.smart2020.org/_assets/files/02_Smart2020Report.pdf.

[3] G. Shen and R. Tucker. *Energy-minimized design for IP over WDM networks*. Journal of Optical Communications and Networking, 1(1):176–186, June 2009. doi:10.1364/JOCN.1.000176.

[4] L. Chiaraviglio, M. Mellia, and F. Neri. *Energy-aware backbone networks: a case study*. In Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on, pages 1–5, Dresden, 2009. IEEE. doi:10.1109/ICCW.2009.5208038.

[5] W. Van Heddeghem and F. Idzikowski. *Equipment power consumption in optical multilayer networks - source data*. Technical Report IBCN-12-001-01, IBCN, University of Ghent, January 2012. Available at: http://powerlib.intec.ugent.be.

[6] R. Hülsermann, M. Gunkel, C. Meusburger, and D. A. Schupke. *Cost modeling and evaluation of capital expenditures in optical multilayer networks*. Journal of Optical Networking, 7(9):814–833, September 2008. doi:10.1364/JON.7.000814.

[7] S. K. Korotky. *Network global expectation model: A statistical formalism for quickly quantifying network needs and costs*. Lightwave Technology, Journal of, 22(3):703–722, 2004. doi:10.1109/JLT.2004.825756.

[8] W. Van Heddeghem, M. De Groote, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester. *Energy-efficiency in telecommunications networks: link-by-link versus end-to-end grooming*. In Optical Network Design and Modeling (ONDM), 2010 14th Conference on, pages 1–6, Kyoto, Japan, 2010. IEEE. doi:10.1109/ONDM.2010.5431570.

[9] C. Dorize, A. Morea, O. Rival, and B. Berde. *An energy-efficient node interface for optical core networks*. In Transparent Optical Networks (ICTON), 2010 12th International Conference on, pages 1–4. IEEE, 2010. doi:10.1109/ICTON.2010.5549142.

[10] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright. *Power awareness in network design and routing*. In Proc. of INFOCOM, Phoenix, USA, pages 1130–1138, April 2008. doi:10.1109/INFOCOM.2008.93.

[11] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker. *Energy Consumption in Optical IP Networks*. Journal of Lightwave Technology, 27(13):2391–2403, July 2009. doi:10.1109/JLT.2008.2010142. Available from: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4815495.

[12] D. C. Kilper, G. Atkinson, S. K. Korotky, S. Goyal, P. Vetter, D. Suvakovic, and O. Blume. *Power Trends in Communication Networks*. IEEE Journal of Selected Topics in Quantum Electronics, 17(2):275–284, March 2011. doi:10.1109/JSTQE.2010.2074187.

[13] F. Idzikowski. *Power Consumption of network elements in IP over WDM networks*. Technical Report TKN-09-006, Technische Universität Berlin, Telecommunication Networks Group, July 2009.

[14] S. Aleksić. *Analysis of Power Consumption in Future High-Capacity Network Nodes*. Journal of Optical Communications and Networking, 1(3):245–258, August 2009. doi:10.1364/JOCN.1.000245.

[15] O. Tamm, C. Hermsmeyer, and A. M. Rush. *Eco-sustainable system and network architectures for future transport networks*. Bell Labs Technical Journal, 14(4):311–327, 2010. doi:10.1002/bltj.20418.

[16] B. Mukherjee. *Optical WDM Networks*. Springer, Reading, Massachusetts, 2006.

[17] C. Belady, A. Rawson, J. Pfleuger, and T. Cader. *Green Grid Data Center Power Efficiency Metrics: PUE and DCIE*. White Paper #6, The Green Grid, 2008.

[18] EPA. *Report to Congress on Server and Data Center Energy Efficiency*. U.S. Environmental Protection Agency, http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf, 2007.

[19] D. W. Martin. *1:1 Protection Switching Overview*. http://www.ieee802.org/1/files/public/docs2007/ay-martin-protection-0707-v01.pdf, July 2007. presentation, San Francisco.

[20] Géant. *The Géant network*. http://www.geant.net/Network/NetworkTopology.

[21] D. Staessens et al. *Value Analysis report on the use of the network planning tool*. Deliverable D7.3, DICONET, European ICT-research project, 2010.

[22] R. Hülsermann, A. Betker, M. Jäger, S. Bodamer, M. Barry, J. Späth, C. Gauger, and M. Köhn. *A set of typical transport network scenarios for network modelling*. ITG FACHBERICHT, (182):65–72, 2004.

[23] H. Zang, J. P. Jue, and B. Mukherjee. *A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks*. Optical Networks Magazine, 1(1):47–60, 2000.

4

A Power Consumption Sensitivity Analysis of Circuit-Switched Versus Packet-Switched Backbone Networks

In this chapter, we refine the results from our 'optical bypass' case studies in Chapter 3. We deepen our earlier analysis by considering the sensitivity of the power saving potential to a wider range of parameters; especially to different artificial and realistic topologies, and transport linerates. The findings on the saving potential of bypassing IP routers are used in Chapter 5.

⋆ ⋆ ⋆

Ward Van Heddeghem, Filip Idzikowski, Francesco Musumeci, Achille Pattavina, Bart Lannoo, Didier Colle, and Mario Pickavet

Accepted for publication in Computer Networks, September 2014

Abstract While telecommunication networks have historically been dominated by a circuit-switched paradigm, the last decades have seen a clear trend towards packet-switched networks. In this paper we evaluate how both paradigms (which have also been referred to as optical bypass and

non-bypass, respectively) perform in optical backbone networks from a power consumption point of view, and whether the general agreement of circuit switching being more power-efficient holds. We consider artificially generated topologies of various sizes, mesh degrees and—not yet previously explored in this context—transport linerates. We cross-validate our findings with a number of realistic topologies.

Our results show that circuit switching is preferable when the average node-to-node demands are higher than half the transport linerates. However, packet switching can become preferable when the traffic demands are lower than half the transport linerate. We find that an increase in the network node count does not consistently increase the energy savings of circuit switching over packet switching, but is heavily influenced by the mesh degree and (to a minor extent) by the average link length. Our results are consistent for uniform traffic demands and realistic traffic demands.

A key take-away message for other research on power saving solutions in backbone networks is that the ratio between the average demand and the demand bitrate has considerable effect on the overall efficiency, and should be taken into account.

4.1 Introduction

Electricity consumption in telecommunication networks is an important issue The worldwide electricity consumption of telecommunication networks (which includes operator networks, office network equipment, and customer premises network access equipment) has been estimated to be 330 TWh in 2012, accounting for 1.7% of the total worldwide electricity consumption in the same year [1]. While it can be argued that this number in itself is relatively small, it is non-negligible and increasing at a rate of 10% per year. Moreover, its relative contribution to the total worldwide electricity consumption is increasing as well (from 1.3% in 2007 to 1.7% in 2012). With the foreseen traffic growth in communication networks [2], this trend is not likely to halt soon. As such, the interest to improve the energy-efficiency of telecommunication networks is a hot research topic, and is of importance for economic (reducing the energy cost), technical (reducing the associated heat dissipation) and environmental (reducing the carbon footprint) reasons.

The electricity consumption in backbone networks is expected to rise considerably The major part of the power consumption in the telecommunication operator networks is currently attributed to the wired aggregation & access networks and mobile radio networks. The backbone networks, in contrast, are estimated to account (in 2012) for only about 8% of the total operator network consumption (which includes the wired aggregation & access, mobile radio and backbone networks) [3]. However, the energy consumption in wired access networks is proportional to the number of connected subscribers, while the consumption in the backbone network is proportional to the traffic volume [3]. With the expected increase of traffic volume, high growth rates in the backbone's energy consumption are expected (potentially even overtaking the access networks' consumption [4]). For this reason, it is important to react timely to the energy issue of backbone networks.

Circuit switching has been identified, so far, as more energy-efficient than packet switching In response, there is a growing body of research literature on reducing the energy consumption in backbone networks. Among the approaches proposed are the introduction of sleep modes, energy-aware routing protocols, energy-aware network design, optical bypassing of power-hungry Internet Protocol (IP) routers, and dynamic rate adaptation. A thorough survey is available in [5]. However, in the last decades, the telecommunication industry has seen a shift from circuit-switched networks to packet-switched networks. There has been some earlier research into the power consumption of circuit switching versus packet switching (briefly discussed in Section 4.2). The general agreement seems to be that circuit switching has a lower power consumption than packet switching.

However, we think that the picture is not so clear-cut Most works point out the benefits of circuit switching over packet switching in terms of power consumption. These benefits depend however on the investigated network scenario. For example, looking at Fig. 4 of [6], the x-axis depicting "Average of random traffic demand" starts from 20 Gbps/node pair, while the capacity of a single Wavelength Division Multiplexing (WDM) channel is set to 40 Gbps. The missing range 0–20 Gbps/node is expected to show that the packet-switched networks can be less power consuming than the circuit-switched networks, as preliminarily indicated in our earlier work [7] and by Bianco et al. in [8].

Contributions of this paper In this paper we extensively compare the circuit and packet-switched IP-over-WDM networks with respect to their power efficiency. We consider circuit switching in the context of optical circuits, in contrast to the more traditional opto-electronic circuit switching such as in SONET/SDH and OTN. We focus on the comparison of circuit switching and packet switching in terms of inverse power efficiency

(W/Gbps), leaving the more complex hybrid solutions aside. The inverse power efficiency is the power (in Watt) required to transport a uniform demand of 1 Gbps (lower values indicate more efficient operation). Note that circuit switching and packet switching in this context has also been referred to as optical bypass (or transparent switching) and non-bypass (or opaque switching) respectively. The four key contributions of our paper with respect to the existing body of research are as follows.

- In addition to considering the mesh degree and network size (in terms of the number of nodes and average physical link length), we evaluate the influence of the channel linerate on the power efficiency of circuit switching versus packet switching, a parameter which to our knowledge has previously not been assessed.
- We particularly look at network scenarios where packet switching is preferable from a power consumption point of view. This aspect has to the best of our knowledge not been addressed in the previous literature (cf. [6], as mentioned above).
- We study the (inverse) power efficiency of both switching paradigms under increasing traffic demand. We show that the power efficiency of packet switching in sparsely-connected networks is almost independent of the traffic demand, whereas for circuit switching the power efficiency improves with increasing traffic.
- We find that a higher node count does not necessarily make circuit switching more preferable. In highly meshed networks the node count does not influence the relative savings of circuit switching over packet switching at all. Our results show that the mesh degree, the demand/linerate ratio and the physical link length are critical parameters.

All in all, our results provide a better insight into the trade-off of the power efficiency of circuit switching versus packet switching.

Organization of this paper We briefly discuss related work in Section 4.2. After outlining the network architecture (Section 4.3) we provide details on our methodology for calculating the network power consumption (Section 4.4). In Section 4.5 we introduce the different set of topologies, traffic matrices and transport linerates that we will consider. Using the result from our dimensioning tool, we show in Section 4.6 that (a) indeed packet switching can be the preferable option with respect to power consumption below certain traffic demand bitrates, (b) that this crossover point is essentially determined by the ratio of the traffic demand over the linerate, and (c) to a minor extent also by the mesh degree.

This paper is an extended version of our earlier work [9]. It includes a more elaborate introduction (Section 4.1) and related work (Section 4.2), a more formal description of our dimensioning algorithm (Section 4.4.2), a validation of our results with demands based on actual traffic measurements from the Abilene topology (Section 4.6.5), an assessment of plausible real-life demand/linerate ratios (Section 4.6.2), a short cross-validation with the results from Shen and Tucker (Section 4.6.4), and a sensitivity analysis to a more detailed IP power consumption model (Section 4.6.6). Moreover, we have now considered 100G linerate technology and dropped the 2.5G linerates (see Table 4.1), and considered different regeneration reaches for the different linerates (see Table 4.1).

4.2 Related work

A decent set of recent papers has focused on the energy-efficiency in optical backbone networks. Some of them have also investigated the differences between the circuit and packet switching paradigms, identified respectively as bypass and non-bypass architectures, in the context of optical networks. In this section, we only focus on the works tackling the case of either establishing a bypass only between a source and target of a traffic demand (circuit switching), or establishing no bypass at all (packet switching).

In [6] Shen and Tucker exploited the concept of lightpath-bypass to perform a power-minimized optical network design, based on Integer Linear Programming (ILP) formulations and heuristics. They distinguish non-bypass (packet switching), direct bypass (circuit switching), as well as an intermediate hybrid solution called multi-hop bypass. A similar problem has been faced in our previous work [7], where simulations and an analytical model were used for the power consumption evaluation of bypass and non-bypass scenarios. In the line of these studies, an analytical model based on expectation values has been also developed by Aleksić and Van Heddeghem in [10], where different variations of the optical bypass strategy are evaluated under different mesh degree scenarios, i.e., from a ring up to full-mesh topologies. Capital Expenditure (CapEx) minimized and power minimized networks designed with an ILP and a genetic algorithm have been considered by Bianco et al. [8]. A bypass and non-bypass architecture (differing by traffic grooming, placement of transponders and (non-)existence of Optical Cross-Connects (OXCs)) in IP-over-WDM are distinguished. Finally, in [11], Aleksić performed a power consumption evaluation of switching and routing elements to compare the circuit and packet switching paradigms, but the analysis is limited to the node level.

Most of these works do not consider the effect of the adoption of different transport linerates on the energy-efficiency of the packet and circuit switching paradigms. In this paper, we extend the above earlier works by analyzing the joint impact of both the mesh degree and different linerates on the energy-efficiency of both switching paradigms. We assess under which conditions each switching paradigm represents the most energy-efficient solution.

4.3 Network architecture: circuit switching versus packet switching

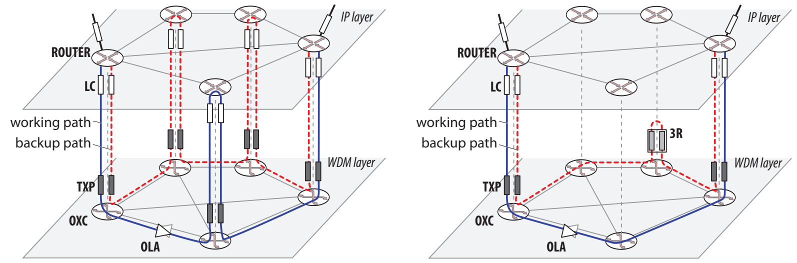

(a) Packet switching **(b)** Circuit switching

Figure 4.1: *The packet-switched and circuit-switched network architectures considered in this paper, showing both the bidirectional working path (solid lines) and backup path (dashed lines) under a 1+1 protection scheme. (LC = Line Card, TXP = Transponder, OXC = Optical Cross-Connect, OLA = Optical Line Amplifier, 3R = 3R regenerator)*

The general architecture of the network is shown in Fig. 4.1 on an example of a 5-node topology (IP/Multiprotocol Label Switching (MPLS) and WDM layers). In the *IP/MPLS layer*, a core router is equipped with line cards, providing one or more ports with short reach interfaces. We assume (differently from [11]) that IP routers have to be present in the backbone network under the circuit switching paradigm, since they exchange the IP traffic with other networks (metro, access) attached to them [12]. The buffers located in the router's line cards are used only at the end nodes of the optical circuits. The granularity of the linerates of the interfaces differs: the access or client-side traffic connects to the router using 1-Gbps interfaces, and the core network side interfaces are either 10-Gbps, 40-Gbps or 100-Gbps interfaces (which we refer to as 10G, 40G and 100G). Note that, depending on the traffic demand bitrate, one or more interfaces can be required per demand.

In the *WDM layer*, long reach transponders with the same capacity as the IP/MPLS layer line cards provide a WDM optical signal, which is switched using an OXC towards the correct physical link. A mux/demux (included in the OXC) aggregates up to 40 channels on a fiber. For each physical link, we assume an unlimited number of fibers to be available. A booster and pre-amplifier (included in the OXC) amplify all channels in a fiber pair respectively upon leaving and entering a node. An Optical Line Amplifier (OLA) is placed every 80 km, and amplifies all channels in a fiber pair. For lightpaths longer than the regenerator span, which depends mainly on the tranponder reach (see Table 4.1), the signal is switched by the OXC to pass through a 3R regenerator.

The way that traffic demands traverse the network is different in packet switching and circuit switching. Under the *packet switching paradigm*, all the traffic in a node—i.e., not only the originating and terminating, but also the transit traffic—is processed at the router in the IP/MPLS layer, as shown by the solid line in Fig. 4.1a. This provides the opportunity to groom traffic, that is bundling traffic belonging to demands from different sources that are destined to the same outgoing link. As a result, the transport channels (wavelengths) are filled more efficiently.

Under the *circuit switching paradigm*, traffic demands traverse the network over a single IP hop, since dedicated optical circuits are set up from the source IP/MPLS node to the target node, as shown by the solid line in Fig. 4.1b. This allows the transit traffic to remain in the optical domain and thus bypass the IP router. For this reason such architectures are often referred to as *optical-bypass* architectures. However, depending on the ratio between the traffic demand bitrates and the channel capacity (i.e., linerate), lightpaths might not be optimally used. For a given set of demands, this might result in a higher number of channels required compared to packet switching.

In both switching cases, we assume a 1+1 protection scheme at the IP layer. Under this scheme, a backup path (dashed line in Fig. 4.1) is simultaneously routed over a link-disjoint physical path with respect to the primary one, so that if the working path fails, the traffic can be instantaneously switched over to the backup path.

4.4 Network dimensioning and power consumption calculation

The intention of our paper is to calculate the power consumption of a set of network topologies given a set of traffic matrices, and this considering both a circuit-switched architecture and a packet-switched architecture. Thereto, we need to specify what power consumption values we consider for the various equipment outlined above, and how we will dimension the network given a certain traffic matrix.

4.4.1 Power consumption model

The power consumption values assumed for each equipment type described earlier in Section 4.3 are listed in Table 4.1. All values are taken from [7] (which goal was to collect and present representative power consumption values for backbone equipment), with the exception of the 40G coherent and 100G coherent transponder values which are based on [13]. The transponder reach, which determines the placement of 3R regenerators, is taken from [14].

Table 4.1: *Power consumption values (source: [7], [13])*

Equipment	Power cons.	Inv. pow. eff.
IP/MPLS 1G-port	10 W	10 W/Gbps
IP/MPLS 10G-port	100 W	10 W/Gbps
IP/MPLS 40G-port	400 W	10 W/Gbps
IP/MPLS 100G-port	1000 W	10 W/Gbps
OLA (per fiber pair, 80 km span)	110 W	-
Transponder 10G non-coherent, reach 3000 km	50 W	5 W/Gbps
Transponder 40G coherent, reach 2500 km	167 W	4 W/Gbps
Transponder 100G coherent, reach 1200 km	389 W	3.9 W/Gbps
3R regenerator xG	2 · transponder xG	-
OXC, 40 ch., with degree d_f	150 W + $d_f \cdot$ 135 W	-

The power-per-port values for the IP router include both the power consumed by the line card and the basic node (i.e., shelves, switch fabric, routing engine, power supply, internal cooling and remaining minor components). We assume the power-per-port value fixed and independent of the load (but not capacity!), as the power consumption of present-day IP routers when idle and under full load are very similar [13, 15]. This also implies that the influence on the power consumption of buffering and table look-up associated with packet switching is negligible.

The power consumption value used for the OXC includes a fixed overhead (150 W) and OXC degree variable part (135 W) that accounts for the switching, mux/demux stages as well as pre- and booster-amplifiers. The OXC degree d_f is defined as the number of network-side bidirectional fiber ports, assuming that all fiber ports are added/dropped at the tributary side (i.e., towards the IP/MPLS layer).

In addition to the total power consumed by the devices listed in Table 4.1, we assume that an equal amount of overhead power is consumed for site cooling and power supply losses, i.e., the Power Usage Effectiveness (PUE) is equal to 2[1].

In Section 4.6.6 we also consider a more accurate IP power model than the above capacity-proportional 10 W/Gbps. In the more accurate power model we account for the actual required IP fabric card shelves, line card shelves, slot cards and port cards. The reason we *do not* use the more accurate model by default is that it introduces some anomalous behavior in the power saving charts, as we will show in Section 4.6.6, thereby somewhat obscuring the general trends.

4.4.2 Dimensioning and power consumption calculation

To calculate and evaluate the power consumption for a given network topology and traffic matrix, for both the packet and circuit-switched architectures, we use a custom Java-based dimensioning tool.

The pseudo-code of the network dimensioning algorithm is given in Alg. 1. The notation used in the description of the network dimensioning method is defined in Table 4.2 with parameters being input to the algorithm, and variables being output of the algorithm. The general steps in dimensioning the network and calculating the power consumption are as follows.

1. First, the traffic D is routed over the physical supply network G with the constraints of the assumed switching paradigm determining also the lightpaths to be established and the number of ports to be installed at the node. For the packet-switched architecture, this means that we consider the logical supply network H (i.e, the IP topology) identical to the physical supply network G, and that traffic over the same physical link $e \in E$ can be groomed. For the circuit-switched architecture, this means that we consider the logical supply network H to be a full-mesh, and no grooming is possible. To achieve 1+1 protection at the IP layer (see Fig. 4.1), the two shortest link-disjoint

---

[1] Note that recently deployed high-capacity data centers with a focus on energy efficiency show much lower PUE values, such as Google claiming to have reached an annualized average PUE across all their tracked data centers of 1.14 by the end of 2011 [16]. However, this is not yet commonplace for telecom operators, with one national operator stating (in private) that '... 1.8–2.0 as an average is not an unreasonable assumption'.

***Table 4.2:** Notation used in the network dimensioning*

	Symbol	Description
Parameters	$G = (V, E)$	directed physical supply network with nodes V and supplied physical links E
	$H = (V, L)$	directed logical supply network with nodes V and supplied logical links L
	C	capacity (bitrate) of a lightpath
	S	length of a span between two OLAs (in kms)
	R	length of a span between two 3R regenerators (in kms)
	W	number of wavelengths per fiber
	D	a Traffic Matrix (TM)
Variables	f_{ij}^{ab}	whether the traffic demand originated at node $a \in V$ and targeted to node $b \in V$ traverses the logical link from $i \in V$ to $j \in V$, $f_{ij}^{ab} \in \{0, 1\}$
	y_l	number of lightpaths established on the logical link $l \in L$, $y_l \in \mathbb{Z}_+$
	z_e	number of fibers installed on the physical link $e \in E$, $z_e \in \mathbb{Z}_+$
	x_i	number of ports (equal to the number of transponders) installed at each node $i \in V$, $x_i \in \mathbb{Z}_+$

physical paths between the source and target nodes are calculated using a minimum cost flow algorithm, where we assume the overall path length, expressed in number of hops, as cost.

2. Then, the wavelength assignment takes place determining also the number of fibers installed at each physical link z_e. This is done in a first-fit fashion [17], meaning that the algorithm finds the first free wavelength/fiber pair that is available on the physical path between source and target nodes.

3. Eventually, the total power of all devices installed in the network is counted using the values from Table 4.1.

    - IP routers and transponders installed at each node $i \in V$ are determined based on the number of ports x_i.

Algorithm 1 Pseudo-code of the network dimensioning and power calculation

```
**Require:** $G, H, C, S, R, W, D$, protectionScheme
**Ensure:** $f_{ij}^{ab}$ for each $(i, j) \in V \times V$ and $(a, b) \in V \times V$, $x_i$ for each $i \in V$, $y_l$ for each $l \in L$, $z_e$ for each $e \in E$
    ($f_{ij}^{ab}$, $y_l$, $x_i$) = routeTraffic($D$, $G$, $H$, $C$, protectionScheme);
    $z_e$ = assignWavelengths($y_l$, $W$, first-fit)
    evaluatePower($x_i$, $y_l$, $z_e$, $S$, $R$);
```

- OXCs installed at each node $i \in V$ are determined based on the number of fibers z_e. Because of the dimensioning tool constraints, we generalized on the OXC power consumption and calculate an average OXC power consumption value based on the average node degree of the network.
- The number of necessary OLAs at each physical link $e \in E$ is determined by its length, length of the span S and the number of installed fibers z_e.
- The number of necessary 3R regenerators at each logical link $l \in L$ is determined by the length of its constituting physical links $e \in E$, and the length of the regenerator span R.

4.5 Case-studies - network scenarios

We will evaluate the power consumption under both a packet-switched and a circuit-switched architecture, and this for a number of different (a) network topologies, (b) traffic matrices, and (c) linerates. This will allow us to do a power consumption sensitivity analysis on an extensive set of parameters.

4.5.1 Topologies

To understand the influence of the connectivity degree and network size (in terms of number of nodes and average physical link length) on the power consumption, we consider a number of *artificially generated topologies*, ranging from minimally meshed (ring) up to maximally meshed (full-mesh) networks, see Table 4.3. For each of these variations we consider networks with the number of nodes N equal to 10, 15, 25, and 33.

To be able to cross-validate our results based on artificial topologies, we also consider four *realistic networks*: the Spanish Telefónica I+D (TID) network model (forecasted potential topology for the year 2020 [18]), the DICONET pan-European Géant network [19], the well-known U.S. NSF network ('us-nobel' at http://sndlib.zib.de/) [20], and the slightly smaller U.S. Abilene network[2] [20]. They are also listed in Table 4.3.

For all of the networks, the IP supply topology H is taken identical to the WDM supply topology G under the packet switching paradigm. All links are bidirectional.

---

[2] The Abilene topology available from sndlib [20] has been slightly modified to represent a survivable network, as required for supporting the 1+1 protection scheme. Thereto, the node ATLAM5 has been removed; all traffic originating from and destined to ATLAM5 has been allocated to the only node it was connected to, i.e., ATLAng.

Table 4.3: *Topologies considered in this study*

Topology	Number of nodes N	Number of bidir. links L	Avg. node degree $\overline{d}$	Mesh degree M	Link length (avg) [km]
ring	10	10	2	0.22	250
ring	15	15	2	0.14	166
ring	25	25	2	0.08	100
ring	33	33	2	0.06	75
half-mesh	10	23	4.5	0.50	250
half-mesh	15	53	7	0.50	166
half-mesh	25	150	12	0.50	100
half-mesh	33	264	16	0.50	75
full-mesh	10	45	9	1.00	250
full-mesh	15	105	14	1.00	166
full-mesh	25	300	24	1.00	100
full-mesh	33	528	32	1.00	75
TID	33	53	3.21	0.10	(52.4)
Géant	34	54	3.18	0.10	(753)
NSF	14	21	3.00	0.23	(1087)
Abilene	11	14	2.55	0.26	(1004)

Similarly to [10] we define the mesh degree M of a network as the ratio of the average node degree of the network under consideration, $\overline{d}$, and the node degree of a full-mesh network having the same number of nodes as the considered network, i.e., d_{mesh} =N-1, so we get $M = \frac{\overline{d}}{d_{mesh}}$. The half-mesh networks have a mesh degree of $M = 0.5$, so the average desired node degree is calculated as $\overline{d} = \frac{N-1}{2}$. To generate these half-mesh networks we (a) start from a ring network with the required number of nodes N and number of links $L_{ring} = N$, (b) then calculate the number of links to add in order to have the desired[3] average mesh (and node) degree, and (c) eventually add these links distributed evenly across the ring (connecting the most-distant nodes, based on the hop count, first). Note that the number of links in such a half-mesh network is given by $L = L_{ring} + N \cdot \frac{\overline{d}-2}{2} = N \cdot \frac{N-1}{4}$.

For the physical link lengths, which influence the power consumption of the OLAs and 3R regenerators, we assume that each of the generated networks covers a geographical area with a diameter of 800 km (which is comparable to a country-sized network such as Germany), or a circumference of approximately 2,500 km. The physical link lengths are then taken

---

[3] Note that, depending on the number of nodes and the requested degree, the theoretical number of links to add might be a fractional number. So we round this value up or down to the closest integer to get a practical (i.e., integral) number of links to add. As a result, the actual degree of the network might differ slightly from the requested one.

to be 2,500 km divided by the number of links in a ring network. For the half-mesh and full-mesh networks we take all other physical links to have the same length, even if this is topologically unrealistic (Table 4.3).

4.5.2 Traffic matrices

For each topology, we generate traffic matrices with *uniform demands*, i.e., an identical demand between each node pair. We consider a range of uniform node-to-node demand values, starting at 1 Gbps, and stepwise increasing up to 220 Gbps. The upper limit of our range is determined so that demands are at least higher than twice our largest considered linerate, which is 100G (see Section 4.3).

Furthermore, for a subset of topologies we also consider more *realistic demand* types. These include random demands, gravity demands, and demands based on actual traffic measurements; similar as for the uniform demands, the demands were scaled to span a large range of actual demands. More details are given in Section 4.6.5, where we perform a sensitivity analysis on the demand type.

4.5.3 Linerates

As noted in Section 4.3 we consider three different transport linerates: 10G, 40G and 100G. This affects the IP interfaces and transponders.

4.6 Results and observations

In this section we compare the power consumption of packet switching (PS) and circuit switching (CS) architectures, evaluated over the artificially generated topologies (from ring to full-mesh) and cross-validated with the realistic topologies[4].

For this evaluation we use three metrics: the absolute total power consumption (kW), the inverse power efficiency (W/Gbps), and the relative power consumption savings of CS over PS (%). The inverse power efficiency is the power (in Watt) required to transport a uniform demand of 1 Gbps (lower values indicate more efficient operation). The relative power consumption savings of CS over PS are calculated as $100 \times \frac{\text{Power}_{\text{PS}} - \text{Power}_{\text{CS}}}{\text{Power}_{\text{PS}}}$, and give a clear indication which switching paradigm is more power-efficient; positive values indicate that CS is preferable, negative values indicate that PS is preferable.

---

[4] Note that a useful extension would be to find the optimum topology (through optimization), instead of comparing given topologies under different conditions. This is however considered out of scope.

Table 4.4: *Overview of our findings. Key findings in bold*

	Finding	Section
General	Sparser topologies consume more	4.6.1
	(Inv.) power efficiency improves with increasing demands, except for PS in sparse topologies	4.6.1
	Higher demands favor CS	4.6.1
d/l ratio	High demand/linerate ratios favor CS, low demand/linerate ratios favor PS	4.6.2
	CS is always preferable for demands higher than half the channel linerate	4.6.2
	There are reasons for real-life networks to operate with demand/linerate ratios (far) above 1	4.6.2
Network size	**Networks with more nodes do not necessarily result in larger relative savings of CS over PS**	4.6.3
	Longer link lengths result in reduced savings for CS	4.6.3
Mesh degree	Savings of CS over PS decrease with increasing mesh degree	4.6.4
	The above behavior is not applicable at low demand/linerate ratios	4.6.4
Demand type	Realistic traffic has a smoother savings profile	4.6.5

As this section is rather dense in content, an overview of the findings in this section is given in Table 4.4, with forward references to the relevant subsections.

4.6.1 General observations

Sparser topologies consume more From Fig. 4.2(a) and (b) we see that sparser topologies (i.e., more ring-like) consume more power than more meshed topologies. This is due to longer paths needed both in the PS and CS.

(Inv.) power efficiency improves with increasing demands, except for PS in sparse topologies Fig. 4.2(c) and (d) show the inverse power efficiency, i.e., the power (in Watt) required to transport a uniform demand of 1 Gbps. We see that the power efficiency of PS (dashed lines) is almost independent of the traffic demand in ring-like networks, whereas in highly-meshed topologies its efficiency gradually improves with increasing traffic. CS (solid lines) behavior is similar to the latter irrespective of the mesh degree.

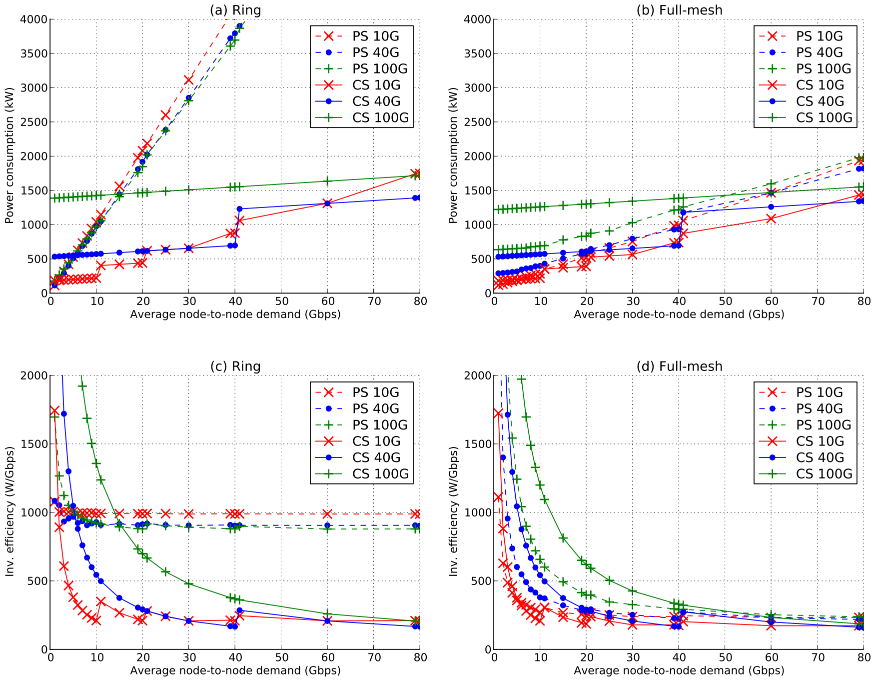

__Figure 4.2:__ The total power consumption and inverse power efficiency of a 15-node ring and full-mesh topology with increasing node-to-node traffic demand. The packet-switched (PS) paradigm shows an overall linear behavior, whereas the circuit-switched (CS) paradigm shows a stepwise behavior whenever the traffic demand becomes a multiple of the channel capacity. The power efficiency of PS in sparsely-connected networks is almost independent of the traffic demand, whereas for CS the power efficiency improves with increasing traffic.

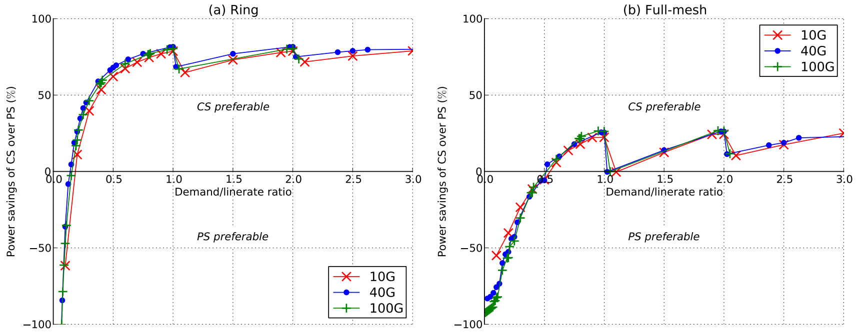

Figure 4.3: Power savings of CS over PS mapped to the ratio of the demand bitrate over the channel linerate (15-node topology). The savings show a stepwise behavior around integral multiples of this ratio (i.e., the savings suddenly drop when the node-to-node traffic demands surpass the channel linerate). The ratio's transition window where CS becomes more preferable than PS is relatively small and relatively independent of the channel linerate (especially for highly-meshed networks, where it is fixed at 1/2).

4.6.2 Influence of the demand/linerate ratio

To get a clear understanding of when CS is more power-efficient than PS (or vice versa), we plot in Fig. 4.3 the power consumption savings of CS over PS. Positive values indicate that CS is preferable, negative values indicate that PS is preferable. For a fair comparison between the different channel linerates, we plot this metric against the ratio of the average demand bitrate over the channel linerate. For a ratio equal to 1, the average demand bitrate is equal to the linerate.

High demand/linerate ratios favor CS, low demand/linerate ratios favor PS Fig. 4.3 shows that increasing demand/linerate ratios lead to higher savings of CS over PS. Low demand/linerate ratios always make PS the preferable paradigm. The reason is that, for low demands, PS can groom traffic into the available capacity of the linerates, whereas for CS low demands result in a lot of unused capacity. Both Fig. 4.3(a) and (b) also clearly show a stepwise behavior around integral multiples of this ratio. This behavior originates from the stepwise behavior of the power consumption of the CS architecture (shown in Fig. 4.2(a)). The CS savings increase until the demand reaches the channel capacity (as there is an increasing usage of the channel capacity), and then suddenly drops when the demands surpass the channel capacity (thereby requiring an extra WDM channel).

CS is always preferable for demands higher than half the channel linerate Fig. 4.3 also indicates that there is a rather narrow transition window of the demand/linerate ratio where CS becomes more preferable than PS. In sparse networks (Fig. 4.3(a)) PS is the preferable option up to about demands being 1/10 to 1/5 of the channel linerate. In highly connected networks (Fig. 4.3(b)), the crossover window is much smaller, and PS is the preferable option for demands being up to half the channel linerate, independently of the utilized transmission technology. The reason that the crossover point is at half the channel linerate is because once a node-to-node demand is larger than half of the channel linerate, there is no free capacity left to groom another demand onto the same channel, and a separate channel is required for each demand.

There are reasons for real-life networks to operate with demand/linerate ratios (far) above 1 Now that we have identified the demand/linerate ratio as an important parameter, the question naturally ensuing from this observation is which demand/linerate ratios are common in real-life networks. Unfortunately, we could not find reliable data on this issue. In [21], Fisher et al. state:

> In backbone networks, pairs of routers are typically connected by multiple physical cables that form one logical bundled link[5] [Doverspike et al., 2010] that participates in the intradomain routing protocol. (...) Link bundles are prevalent because when capacity is upgraded, new links are added alongside the existing ones, rather than replacing the existing equipment with a higher capacity link. (...) Bundled links are also necessary when the aggregate capacity of the bundle exceeds the capacity of the fastest available link technology. In today's backbone networks, a vast majority of links would be bundled, with bundles consisting of two to approximately twenty cables, a majority between the two extremes.

Note that link bundles of 'two to twenty cables' (with 'cables' corresponding in this context to wavelengths or channels) imply a demand/linerate ratio of 2–20 as well. The range for this ratio in Fig. 4.3 (and later figures) is only up to 3. The referenced work (Doverspike et al. [22]) does mention a third driver for link bundling, which is resilience and consequently network stability. If one of the component links fails, the bundled link remains

---

[5] Link bundling is also referred to under various other umbrella terms such as link aggregation, link bonding, link teaming and port trunking. The IEEE 802.1AX-2008 standard uses the term 'link aggregation'.

up and a failure-driven topology update is not required. Unfortunately, [22] does not provide actual data (such as link bundle counts for real operators) to ground their—otherwise plausible—claims. In [23], the availability of bundled links (referred to as 'parallel paths') are an important premises for one of the proposed energy-saving solution, but no actual data or references are given that give insight to what extent this is actually the case in current backbone networks. On the contrary, the work admits that splitting IP traffic demands over multiple parallel paths "is a strong assumption, as multi-path routing is normally not enabled in today's routers. MPLS allows this kind of traffic engineering, but the label switched paths (LSP) are not frequently reconfigured today, either." In an expert interview with a large national operator, we were informed that the choice of implementing a link through either multiple smaller capacity interfaces or through a single overprovisioned interface is largely governed by economic (i.e., cost) decision. Both options were feasible, i.e., operation with a demand/linerate ratio above as well as below 1 exists. Actual data was unfortunately not available. Summarized, the information above suggests that there are some good reasons for real-life networks to operate with demand/linerate ratios (far) above 1, and that this is at least done in some cases. This would imply that a CS architecture is more desirable than a PS architecture from a power consumption point of view.

4.6.3 Influence of the network size (number of nodes and physical link lengths)

Fig. 4.4 shows the power consumption savings of CS over PS for networks with different number of nodes (the network with N=25 has been omitted for clarity). The subfigures (a) to (d) correspond to an increasing mesh degree. Fig. 4.4(b) represents a mesh degree M=0.1, and contains in addition two realistic topologies that also have M=0.1 (the lowest mesh degree of the 10-node and 15-node topology is higher than 0.1, see Table 4.3).

Networks with more nodes do *not* necessarily result in larger relative savings of CS over PS For sparse topologies (Fig. 4.4(a) through (c)) the node count has considerable influence on the relative savings of CS over PS. For the ring topology, a higher node count makes CS more preferable. This is due to the higher hop count in larger ring networks, which implies a much higher IP-layer contribution, which increases the PS power consumption. This is inline with [6]. However, our results indicate that the above rule cannot be applied universally to all sparse topologies. In Fig. 4.4(c) a higher node count does not consistently correspond to increased CS sav-

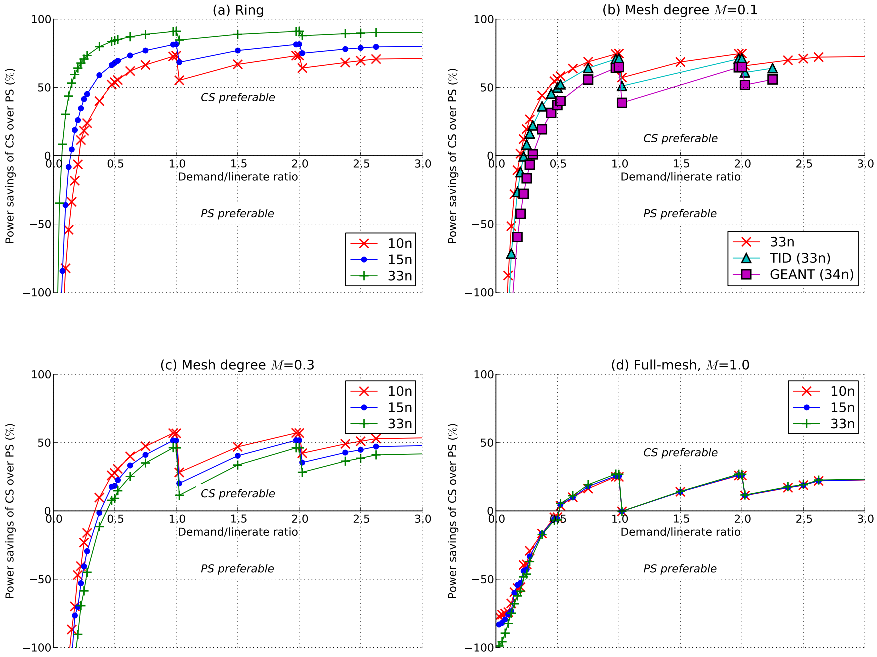

Figure 4.4: *Influence of the node count on the power savings of CS over PS (for linerate = 40G). Only for sparse topologies (i.e., (a) through (c)) the node count has an influence on the savings. While for a ring topology a higher node count leads to more savings, this is not consistently the case for other sparsely meshed topologies. The relatively large deviation of the Géant topology from the general trend is explained in Fig. 4.5.*

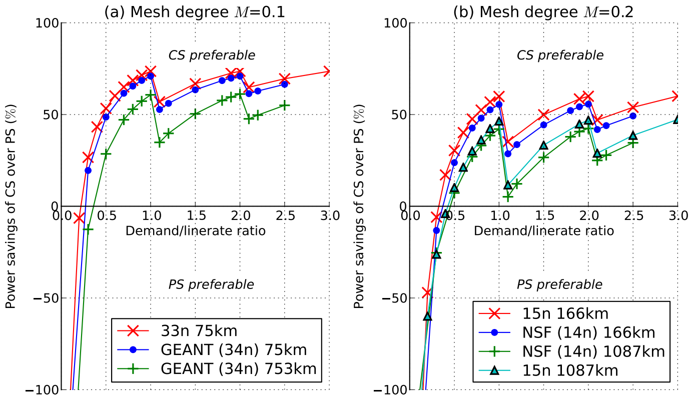

__Figure 4.5:__ Influence of the average physical link length on the relative savings of CS over PS (linerate = 10G). Longer link lengths result in lower savings, and explain why the savings profile of topologies such as Géant (average physical link length = 753 km) does not correspond very well with our artificial topology of the same node count but much shorter link length.

ings (the savings for 33-node artificial topology are lower than for the 15-node topology). Moreover, while in Fig. 4.4(b) the realistic TID network (33 nodes) savings seems to be inline with the 33-node artificial topology, the Géant network (34 nodes) curve is considerably lower. There must be another parameter with substantial influence on power savings.

Longer link lengths result in reduced savings for CS In order to explore the reason of the above described anomaly, Fig. 4.5(a) plots, in addition to the 33-node artificial topology (physical link length = 75 km) and the original Géant topology (average physical link length = 753 km), the same Géant topology where all links have been (artificially) set to 75 km. The figure shows that the difference in link length is the reason of the diverging behavior of the original Géant topology from the artificial 33-node topology. The long link length of the original Géant topology increases the number of required OLAs and 3R regenerators and the associated power consumption. As the additional power consumption has a larger relative impact on the CS power consumption, the power consumption savings of CS over PS decrease accordingly. This is also confirmed by Fig. 4.5(b) where the NSF network (14 nodes, mesh degree M = 0.2, average physical link length = 1087 km) is compared with our artificial 15-node M = 0.2 topology. When the link lengths are adjusted (either from the artificial topology, or from the NSF network), the savings curves become very similar.

4.6.4 Influence of the mesh degree

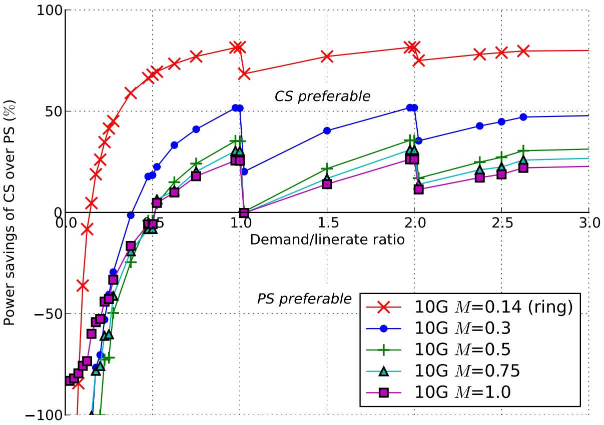

Figure 4.6: Influence of the mesh degree on the relative savings of CS over PS (for linerate = 40G, and 15-node topologies). Higher mesh degrees (M) result in lower savings.

Although we have not focused on the mesh degree yet, it is already clear from the previous figures and discussion that this parameter is of considerable influence on the power savings of CS over PS.

Savings of CS over PS decrease with increasing mesh degree As shown in Fig. 4.6, the savings of CS over PS tend to decrease for increasing mesh degree, as adding more edges decreases the hop count and thus more interfaces (i.e., router ports and transponders) can be saved in intermediate nodes of the PS architecture while still performing traffic grooming. On the other hand, for the CS architecture, a higher mesh degree only impacts the OLAs (and eventually, the regenerators) consumption, which constitutes a less relevant contribution in the total consumed power if compared to the power spent by the interfaces. It is useful to remark that this is also confirmed by the data in the study by Shen and Tucker [6]: if we calculate the mesh degree for the three networks considered in [6], then indeed higher mesh degrees correspond to reduced preference for CS (see Table 4.5).

The above behavior is not applicable at low demand/linerate ratios An exception to this behavior is obtained for low demand/linerate ratios (i.e., higher channel linerates under low traffic conditions). This is shown in Fig. 4.7, which plots the savings as a function of the mesh degree for different demand/linerate ratios. In this case, passing from ring to half-mesh

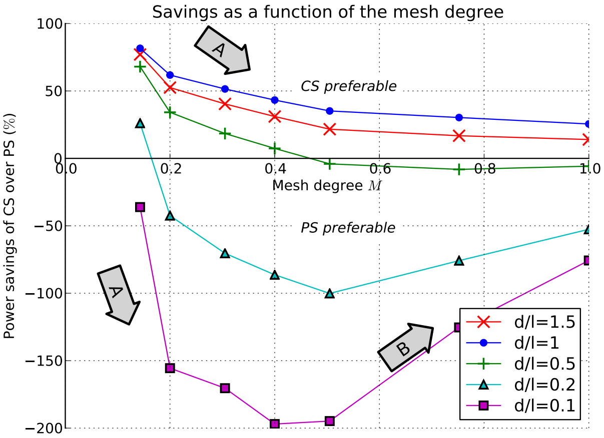

Figure 4.7: *Influence of the mesh degree and the demand/linerate ratio d/l on the relative savings of CS over PS (for linerate = 40G, and 15-node topologies). For low demand/linerate ratios there is an optimum point where PS is favorable.* Effect A: *The PS power consumption decreases with increasing mesh degree, as less IP hops are required.* Effect B: *The PS power consumption increases towards that of the CS solution, because the PS grooming potential decreases with increasing mesh degree.*

topologies has, as previously, a higher benefit for the PS than for the CS solution (i.e., the PS solution benefits from the reduction in IP hops, shown as Effect A in Fig. 4.7). However, adding further links to the network (i.e., going towards full-mesh topologies), does *not* require more interfaces for the CS solution, but it does for the PS solution. However, there is lower opportunity for traffic grooming, so with high channel linerates interfaces are underutilized, thus causing higher relative power consumption of the PS solution compared to the CS solution (shown as Effect B in Fig. 4.7).

It is interesting to point out that for the full-mesh case the power consumption of PS and CS are *not* equal (i.e., CS over PS savings are not zero), as one might incorrectly expect. The link disjoint backup paths always require two hops in both switching paradigms, but the intermediate node

Table 4.5: *Calculating the mesh degree for the 3 networks considered in [6] confirms the finding that higher mesh degrees correspond to reduced preference for CS*

Topology [6]	Reported CS over PS saving [6]	Mesh degree
6 nodes, 8 links	25% savings	0.53
15 nodes, 21 links	40% savings	0.20
42 nodes, 43 links	45% savings	0.09

requires IP ports under the PS paradigm only, leading to CS being more preferable. However, an exception to this is observed for high linerates (e.g., 100G), combined with low demands bitrate (e.g., 15 Gbps per demand). In this case the opportunity to groom traffic in the PS scenario produces higher power benefits in comparison to the high demands bitrate situation, and thus the CS option is outperformed.

To summarize our finding for the mesh degree: in general, the power consumption advantage of CS over PS decreases with an increasing mesh degree. In other words, for networks with a lower mesh degree (such as ring networks), CS is *more* preferable than it is for fully meshed networks. However, this observation does not necessarily hold for low traffic conditions (i.e., demand/linerate ratio < 0.5) as in that case the effect of underutilized channels in the CS scenario starts to dominate.

4.6.5 Sensitivity to non-uniform demands

In all of the above scenarios we assumed fully-meshed uniform demands. To see the effect of non-uniform demands on the power savings of CS over PS, we consider in Fig. 4.8 two additional demand types: (a) a gravity traffic matrix where nearby nodes have larger demands, thus closer resembling real life demands [19], and (b) a random fully meshed traffic matrix where each demand is evenly distributed between -30% and +30% of the nominal demand.

We consider both demand types for two of our network topologies where we have realistic demands available. For the Géant topology the gravity traffic matrix is based on a mathematical model [19]. For the Abilene topology, the traffic matrix is based on actually measured traffic, as made available through sndlib [20]. It was generated by taking the maximum for each demand over the period 2004-07-01 to 2004-07-31 (per [24]), scaled up to a set of traffic matrices with the average demand bitrate ranging from 1 Gbps to 25 Gbps. As we already noted in Section 4.5.1, we have removed the ATLAM5 node and moved all its traffic to the ATLAng in order to have a survivable network that can provide 1+1 protection for each traffic demand. Furthermore, since the Abilene traffic data is given as unidirectional traffic, and we only consider bidirectional demands, we have taken the maximum value of node-to-node demands where demands in each direction were different. Finally, we converted the original Mbps demands to Gbps, rounding up to the nearest integer value.

Realistic traffic has a smoother savings profile While the saving curves associated with uniform demands show the distinct stepwise behavior, the

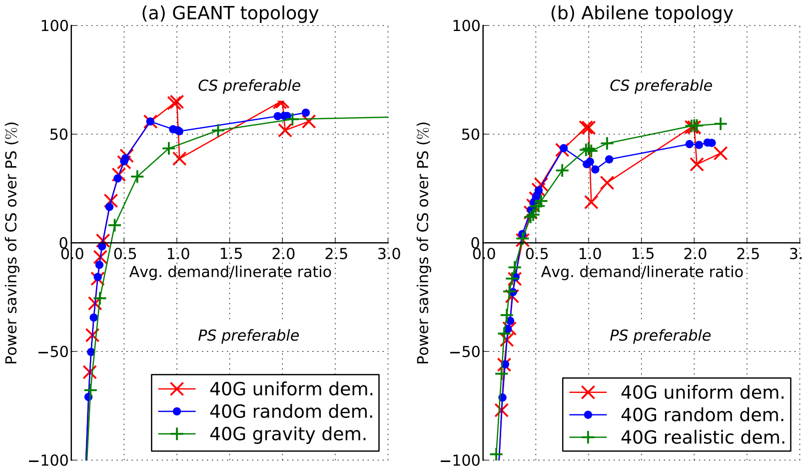

Figure 4.8: *Influence of different demand types on the savings of CS over PS, both for the (a) Géant and (b) Abilene topology (linerate=40G). While uniform demands show a distinct stepwise behavior, more realistic demand sets (i.e., random and gravity demands) smooth out this behavior.*

curve is much smoother for random demands and gravity demands. This stepwise behavior originates in the stepwise power profile of the CS architecture, as explained in Section 4.6.2 and shown in Fig. 4.2(a). For more realistic traffic, the network's *average* demand/linerate ratio is the result of a mix of different demand values (the demand between each node pair is potentially different). As such the behavior that occurs when a demand is just below or just above the linerate is smoothed out. However, the general trend observed before remains valid: CS is preferable for demands higher than half the channel linerate (on average) also under the gravity and random traffic matrices. This observation holds for both the Géant gravity demands (generated based on a mathematical model), as well as the Abilene realistic demands (which are based on real traffic measurements in the Abilene network). This increases our confidence that our results hold for other real demands as well.

4.6.6 Sensitivity to a more detailed IP power consumption calculation

In Section 4.4.1 we modelled the power consumption of IP backbone routers as 10 W/Gbps. The implication of this is that we assume perfect power-proportionality of backbone routers to changes in the required capacity.

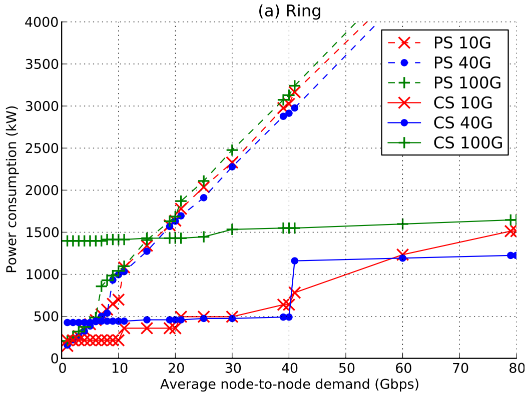

Figure 4.9: *Influence of using more accurate IP power consumption modelling on the total power consumption (15-node topology). The overall shape of the curves is not affected; compare with Fig. 4.2(a).*

This is a simplification, as IP backbone routers consist of various building blocks, and a slight increase in capacity might require the addition of a certain component, incurring a significant additional power consumption. For example, Cisco's CRS series consists of fabric card shelves, line card shelves, slot cards and port cards (see [25] for a concise overview).

To assess the impact and validity of our 10 W/Gbps simplification, we have also calculated the power consumption and power consumption savings of CS over PS with a more accurate calculation. For each IP node, we determine the required number of fabric card shelves, line card shelves, slot cards and port cards, and multiply this with the associated equipment power consumption values[6].

The result of the more detailed calculation is shown in Fig. 4.9 and Fig. 4.10.

We make the following two observations:

- As expected, Fig. 4.9 shows that the power consumption increases in a more stepwise fashion than in Fig. 4.2(a). For example, for the PS 100G case the increase from an average node-to-node demand of 6 Gbps to 7 Gbps requires in each node a second line card shelf and subsequently a fabric card shelf to connect both line card shelves, which is clearly visible[7].

---

[6] We have used the capacity and power consumption data from Table 1 in [7], with the power consumption as follows: fabric card shelf=8100 W, line card shelf=2401 W, 140G slot card=401 W, 14×10G port card=135 W, 3×40G port card=315 W, 1×100G port card=135 W.

[7] For the 15-node network, including a PUE=2, the fabric card shelf (8.1 kW) and the line card shelf (2.4 kW) account for $2 \times 15 \times (8.1 + 2.4) = 315$ kW in the sudden jump.

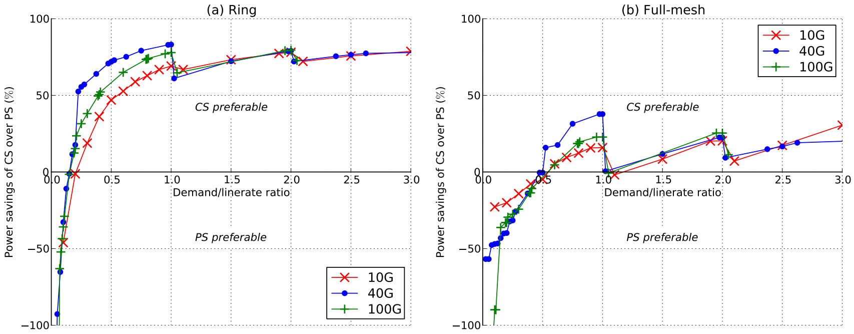

Figure 4.10: Influence of using more accurate IP power consumption modelling on the power savings of CS over PS (15-node topology). The overall shape of the curves is not affected; compare with Fig. 4.3.

- The overall shape of the CS over PS savings curves are not affected, as can be seen when comparing Fig. 4.10 with Fig. 4.3. While specific capacity requirements might result in unfortunate combinations when comparing CS and PS—such as can be seen in Fig. 4.10(b) for the 10G architecture at a demand/linerate ratio of 1.1, where the CS architecture suddenly requires an additional fabric card shelve in each of the 15 nodes, while the PS architecture does not—the overall shape of the curve is nearly identical to those which we calculated using the power-proportional 10 W/Gbps approach.

Thus, the general trends and conclusion from the earlier sections are not affected when using the more accurate IP power consumption calculation. On the downside, the more accurate IP power modeling does introduce sudden changes in the shape of the curves that can only be properly explained when looking in detail at the resulting data. Furthermore, the location of these 'jumps' does not necessarily correspond to real-life deployments, as actual equipment deployment depends also on such issues as expected traffic growth [13]. For these reasons, we eventually chose to present our results with the more idealized 10 W/Gbps model, as it more clearly shows the overall behavior of using a CS versus PS architecture.

4.7 Conclusion and further work

In this paper we extensively compared the power consumption of circuit and packet switching architectures in optical backbone networks. We evaluated the impact of the channel linerate, the network size (both number of nodes and physical link length), demand/linerate ratio and the network mesh degree to assess under which conditions each switching paradigm represents the most power-efficient solution.

We found that, in general, circuit switching is preferable, as fewer power-hungry IP router ports and WDM transponders are needed. This is especially true for networks that use link bundling, i.e., where a node-to-node logical link with a certain capacity is realized using multiple links/interfaces with smaller capacities. We did not find unambiguous data on how common or uncommon link bundling is in backbone networks, but the few sources we were able to find suggest that it is attractive from an operational point of view, and used in at least some cases. However we point out on the top of the related work that for relatively low traffic values—i.e., when the demands bitrate is lower than at least half the channel linerate—the packet switching solution is more power-efficient, thanks to the opportunity of exploiting traffic grooming to better utilize network resources.

Our key finding is that an increase in the network node count does not consistently increase the power savings of circuit switching over packet switching. Instead, these power savings are heavily influenced by the mesh degree and (to a minor extent) by the average physical link length. Increasing the network mesh degree produces higher energy benefits for packet switching than for circuit switching, as more power can be saved in intermediate nodes in the former case. Our main analysis was performed using uniform traffic demands, however we cross-validated our results using more realistic demand sets and found that the key results hold. In fact, more realistic demands remove erratic behavior from the power savings of circuit switching over packet switching that is otherwise observed when the average node-to-node demand bitrate is slightly higher than the transport linerate.

A message to take away by researchers looking into power saving solutions in backbone networks, is that the assumptions made with respect to the average demand bitrate and transport linerates do matter. For example, evaluating a particular solution in a scenario with an average demand of 20 Gbps over either 10G interfaces or 40G interfaces affects the overall network power efficiency beyond just the slight increase in power efficiency associated with 40G interfaces. It is important to be aware of this for a fair evaluation.

While our work already evaluated the circuit switching vs. packet switching paradigm over a wide set of input parameters, there are still a number of interesting inputs to consider for useful further work. First, while our power model accounts for the required equipment capacity, we did assume that the power consumption does not vary for different traffic loads, which is a reasonable assumption for current backbone equipment. However, should the power consumption of future packet switches become more proportional to the load, it is likely that this will influence the outcome of our comparison. Furthermore, we assumed that at each node there is both an IP switch and a WDM switch; this is not always the case in backbone networks, which motivates further investigation of such heterogeneous networks. Finally, our study focussed on two extreme scenarios, being either packet switching or circuit switching. Hybrid solutions—e.g. those that perform a joint optimization of the IP and WDM layer (such as the multi hop bypass solution in [6]), the use of multi-linerate transponders, or hybrid forms of packet and circuit switching—would likely perform optimally under a wider range of traffic/demand linerate ratios. Because of the intention of the current study, this was out of scope. It would however be a useful research topic to compare the impact and optimization potential of such hybrid solutions.

Acknowledgments

The work described in this paper was carried out with the support of the iMinds-project GreenICT and the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257740 (Network of Excellence TREND).

References

[1] W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester. *Trends in Worldwide ICT Electricity Consumption from 2007 to 2012*. Computer Communications, 50:64–76, September 2014. doi:10.1016/j.comcom.2014.02.008.

[2] Cisco. *Cisco Visual Networking Index: Forecast and Methodology, 2010-2015*, 2011.

[3] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch. *Energy consumption of telecommunication networks and related improvement options*. Journal of Selected Topics in Quantum Electronics, 17(2):285–295, 2011. doi:10.1109/JSTQE.2010.2053522.

[4] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker. *Energy Consumption in Optical IP Networks*. Journal of Lightwave Technology, 27(13):2391–2403, July 2009. doi:10.1109/JLT.2008.2010142.

[5] A. P. Bianzino, C. Chaudet, D. Rossi, and J.-L. Rougier. *A survey of Green Networking Research*. IEEE Commun. Surv. Tutorials, 14(1):3–20, 2012. doi:10.1109/SURV.2011.113010.00106.

[6] G. Shen and R. Tucker. *Energy-minimized design for IP over WDM networks*. Journal of Optical Communications and Networking, 1(1):176–186, June 2009. doi:10.1364/JOCN.1.000176.

[7] W. Van Heddeghem, F. Idzikowski, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester. *Power consumption modeling in optical multilayer networks*. Photonic Network Communications, 24:86–102, October 2012. doi:10.1007/s11107-011-0370-7.

[8] A. Bianco, E. Bonetto, F. Musumeci, A. Pattavina, and M. Tornatore. *CapEx/OpEx Evaluation of Circuit vs Packet Switched Optical Networks*. In Proc. of ONDM, Brest, France, pages 5973–5977, April 2013.

[9] W. Van Heddeghem, F. Musumeci, F. Idzikowski, A. Pattavina, B. Lannoo, D. Colle, and M. Pickavet. *Power Consumption Evaluation of Circuit-Switched Versus Packet-Switched Optical Backbone Networks*. In Proc. of IEEE OnlineGreenComm, online, 2013.

[10] S. Aleksić, W. Van Heddeghem, and M. Pickavet. *Scalability and Power Consumption of Static Optical Core Networks*. In Proc. of Globecom, Anaheim, USA, pages 3465–3471, December 2012. doi:10.1109/GLOCOM.2012.6503651.

[11] S. Aleksić. *Analysis of Power Consumption in Future High-Capacity Network Nodes*. Journal of Optical Communications and Networking, 1(3):245–258, August 2009. doi:10.1364/JOCN.1.000245.

[12] X. Dong, T. El-Gorashi, and J. Elmirghani. *On the Energy Efficiency of Physical Topology Design for IP Over WDM Networks*. Journal of Lightwave Technology, 30(12):1931–1942, June 2012. doi:10.1364/JOCN.1.000176.

[13] W. Van Heddeghem, F. Idzikowski, E. Le Rouzic, J. Mazeas, H. Poignant, S. Salaun, B. Lannoo, and D. Colle. *Evaluation of Power Rating of Core Network Equipment in Practical Deployments*. In Proc. of IEEE GreenCom, online, pages 126–132, September 2012. doi:10.1109/GreenCom.2012.6519628.

[14] G. Rizzelli, A. Morea, M. Tornatore, and O. Rival. *Energy Efficient Traffic-Aware Design of On–Off Multi-Layer Translucent Optical Networks*. Computer Networks, 56(10):2443–2455, July 2012. doi:10.1016/j.comnet.2012.03.013.

[15] A. Vishwanath, J. Zhu, K. Hinton, R. Ayre, and R. S. Tucker. *Estimating the Energy Consumption for Packet Processing, Storage and Switching in Optical-IP Routers*. In Proc. of OFC, Anaheim, USA, 2013. doi:10.1364/OFC.2013.OM3A.6.

[16] Google. *Data Center Efficiency – Power Usage Effectiveness*. http://www.google.com/about/datacenters/inside/efficiency/power-usage.html, 2012.

[17] H. Zang, J. P. Jue, and B. Mukherjee. *A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks*. Optical Networks Magazine, 1(1):47–60, 2000.

[18] F. Idzikowski et al. *Final report for the IRA "Energy-efficient use of network core resources*. Deliverable D3.3, TREND Project, 2012.

[19] D. Staessens et al. *Value Analysis report on the use of the network planning tool*. Deliverable D7.3, DICONET, European ICT-research project, 2010.

[20] S. Orlowski, M. Pióro, A. Tomaszewski, and R. Wessäly. *SNDlib 1.0–Survivable Network Design Library*. Networks, 55(3):276–286, 2010. http://sndlib.zib.de/.

[21] W. Fisher, M. Suchara, and J. Rexford. *Greening backbone networks: reducing energy consumption by shutting off cables in bundled links*. In

Proc. of ACM SIGCOMM workshop on Green networking, New Delhi, India, pages 29–34, 2010. doi:10.1145/1851290.1851297.

[22] R. Doverspike, K. K. Ramakrishnan, and C. Chase. *Structural overview of ISP networks*. In C. Kalmanek, S. Misra, and R. Yang, editors, Guide to Reliable Internet Services and Applications, pages 19–93. Springer, 2010. doi:10.1007/978-1-84882-828-5_2.

[23] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz. *Dynamic routing at different layers in IP-over-WDM networks - Maximizing energy savings*. Optical Switching and Networking, 8(3):181–200, 2011. doi:10.1016/j.osn.2011.03.007.

[24] E. Bonetto, L. Chiaraviglio, F. Idzikowski, and E. Le Rouzic. *Algorithms for the Multi-Period Power-Aware Logical Topology Design with Reconfiguration Costs*. Journal of Optical Communications and Networking, 5(5):394–410, May 2013. doi:10.1364/JOCN.5.000394.

[25] W. Van Heddeghem and F. Idzikowski. *Equipment power consumption in optical multilayer networks - source data*. Technical Report IBCN-12-001-01, IBCN, University of Ghent, January 2012. Available at: http://powerlib.intec.ugent.be.

5

A Quantitative Survey of the Power Saving Potential in IP-over-WDM Backbone Networks

In this chapter, we use the analytical power model from Chapter 3 to quantitatively assess the total impact of various power saving techniques in backbone networks. Bypassing IP routers as we studied it in Chapter 4 is one such technique; we use our findings from that chapter as an input.
We'd like to point out that the notation of the analytical model has been changed and optimized slightly in comparison to the one presented in Chapter 3.

⋆⋆⋆

Ward Van Heddeghem, Bart Lannoo, Didier Colle, Mario Pickavet, and Piet Demeester

Submitted to IEEE Communications Surveys and Tutorials, May 2014

Abstract The power consumption in Information and Communication Technology (ICT) networks is growing year by year; this growth presents challenges from a technical, economic and environmental point of view. This has lead to a great number of research publications on 'green' telecommunication networks. In response, a number of survey works have appeared

as well. However, with respect to backbone networks most survey works (a) do not allow for an easy cross-validation of the savings reported in the various works, (b) nor do they provide a clear overview of the individual and combined power saving potential.
Therefore, in this work we survey the reported saving potential in IP-over-WDM backbone telecommunication networks across the existing body of research in that area. We do this by mapping more than 10 different approaches to a concise analytical model, which allows us to estimate the combined power reduction potential.
Our estimates indicate that the power reduction potential of the once-only approaches is 2.3× in a Moderate Effort scenario and 31× in a Best Effort scenario. Factoring in the historic and projected yearly efficiency improvements (“Moore’s law”) roughly doubles both values on a 10-year horizon. The large difference between the outcome of the Moderate Effort and Best Effort scenario is explained by the disparity and lack of clarity of the reported saving results, and by our (partly) subjective assessment of the feasibility of the proposed approaches. The Moderate Effort scenario will not be sufficient to counter the projected traffic growth, although the Best Effort scenario indicates that sufficient potential is likely available. The largest isolated power reduction potential is available in improving the power associated with cooling and power provisioning, and applying sleep modes to overdimensioned equipment.

5.1 Introduction

Power consumption in backbone telecommunication networks is still growing The global amount of Internet Protocol (IP) traffic is growing every year. While this growth is gradually slowing down from an earlier Compound Annual Growth Rate (CAGR) of 100% (about 10 years ago) to an estimated CAGR of around 20–30% currently, this reduced growth still outperforms the annual 13% efficiency increase of new telecommunication equipment in the backbone network [1]. As can be seen in Fig. 5.1 this creates a so-called ‘energy gap’, and as such, the power consumed by telecom backbone network devices continues to increase year by year. This presents issues both from an economic (reducing the energy cost), technical (reducing the associated heat dissipation) and environmental (reducing the carbon footprint) point of view. And while today the power consumption in backbone networks makes up only about 8% of the total operator network power consumption (which includes the wired access, mobile access and backbone network) [2], with the expected increase of traffic volume, high growth rates in the backbone’s energy consumption are expected, po-

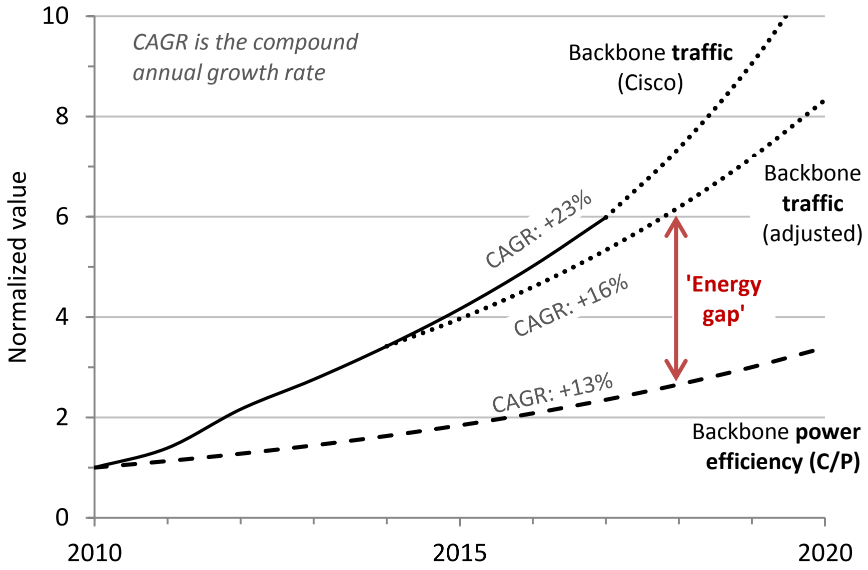

Figure 5.1: Traffic growth and new equipment energy efficiency improvement in the IP backbone, normalized to the year 2010. As the traffic is growing faster than the efficiency improvement, this creates an 'energy gap'. Traffic estimations based on Cisco forecasts as detailed in Section 5.4.3, backbone efficiency ([C/P]=Gbps/W) improvement at +13% per year from Section 5.4.4.

tentially even overtaking the access network's consumption[1] [3]. The three issues mentioned above—economic, technical and environmental—are reflected in the increasing number of publications and research on this 'green' networking topic by academia, industry and governmental bodies alike [4].

Most 'classical' surveys of power saving approaches do not list and quantify the power saving potential The research on 'green' ICT networks can be categorized in two broad categories: (a) estimating (current and future) network power consumption on the one hand; and (b) proposing and evaluating novel techniques to reduce the power consumption on the other hand. Especially for the latter, there are a number of surveys available that list and categorize power saving approaches. Notable ones focusing (partly) on optical backbone networks include those from Zhang et al. [5], Bolla et al. [6] and Bianzino et al. [7] as listed in Table 5.1. We discuss these works in slightly more detail in Section 5.2. Categorization of the power

---

[1] The reason is that the power consumption in wired access networks is proportional to the number of connected subscribers, while the consumption in the backbone network is proportional to the traffic volume [2].

saving approaches in the different survey works is done using a variety of different criteria, such as the area of application (e.g. circuit-level versus network-level) or the considered time scale (e.g. sleep mode during operation versus energy-efficient network design in the planning phase). While these surveys certainly are worthwhile to make sense of the growing body of publications on this topic, it is a striking observation that none of the above three works list the power saving potential reported by (most of) the works they survey in a way that allows an easy cross-comparison. This is curious, as it would seem that the estimated (order of) energy saving potential is one of the main relevant outcomes in most of the surveyed works. Furthermore, such an inventory would give an idea of the relative power saving potential of the different approaches, and the consensus across the research community. This would allow—at least as a first step—to assess which techniques are most promising and should be prioritized from an energy saving point of view; a complete assessment would require consideration of other aspects such as implementation cost and associated operational issues.

Table 5.1: *An overview of surveys on energy efficiency in (backbone) telecommunication networks*

Publication	Scope	Reported power savings
Zhang, 2010 [5]	Optical networks (fixed access, metro, core)	-
Bolla, 2011 [6]	Wired networks	-
Bianzino, 2012 [7]	Wired networks	-
Roy, 2008 [8]	Wireline and wireless	Energy saving potential: wireline ≈ 40%, wireless ≈ 60%
Lange, 2011 [2]	Home networks up to backbone networks	Energy saving potential in 2007: backbone 65%, total network 56%
Parker, 2011 [9]	Photonic telecom networks	Energy efficiency improvement: 1000×
Kilper, 2012 [4]	Optical transmission networks	-
Green Meter, 2013 [10]	Mobile access, wireline access, core network	Energy efficiency improvement (2010–2020): mobile access 1043×, wireline access 449×, core 64×
Le Rouzic, 2013 [11]	Optical networks (fixed access, metro, core, data centers)	Individual savings only, no consolidated savings
Le Rouzic, 2013 [12]	Metro, core, data centers	Individual savings only, no consolidated savings

Existing works that do survey the (combined) power saving potential have some shortcomings To be fair, a number of works exist that *do* map the power saving potential of different power saving techniques, and/or their combined power saving potential. We have listed them in Table 5.1. However, we think each of these works suffers from a number of different shortcomings. Specifically, we found that some studies either (a) lack sufficient references or traceability of the quoted savings ([2, 8]), (b) provide an assessment that is too rough ([9]), or (c) do not provide a combined saving potential ([11, 12]). Finally, while GreenTouch's Green Meter [10] does address the above issues, the limited number of referred works does not help the reader to get an idea on the agreement of the given power savings across the research community outside of GreenTouch. We discus the above cited works in more detail in Section 5.2.

We address the above shortcomings based on an analytical model to survey and quantify the power saving potential in backbone networks Our earlier work [13] was a first response to the various issues identified above; we complete that initial—but incomplete, as it did not contain a consistent and methodological survey, and no saving potential was calculated—proposal in this study. Our methodology is based on a concise analytical model to estimate the power consumption of a backbone network[2]. A comparable model was first used by Baliga et al. [3] and later formalized by Kilper et al. [1]. In [16], we have used a similar analytical formulation in our study on representative power consumption values for backbone equipment. It was not used to survey the power saving potential however. In essence, the total network power consumption is estimated by multiplying the power rating (W/Gbps) of all relevant backbone network equipment with the estimated total amount of traffic in the network. A number of correction factors are introduced in the above multiplication to account for the impact of traffic protection, overprovisioning, cooling power consumption, etc. In this study, we survey and map different energy saving approaches to each of the factors in the analytical model. This has the advantage that the interdependence of the various power saving approaches can be reduced or assessed more easily. Also, the model allows us to apply the impact of specific power saving approaches to specific equipment (such as only to IP routers, and not to transponders). Finally, as the model is based on a set of simple multiplications, we can rather easily calculate the combined impact

---

[2] An alternative technique to estimate the network power consumption is to use a top-down approach, where the total telecom operator power consumption is trimmed from non-relevant contributors (such as data centers, telecom offices) to give only the network relevant contribution. This has been used by e.g. Malmodin et al. [14] and Lambert et al. [15] for estimating the worldwide telecom operator network electricity consumption.

of all power saving approaches.

Key contributions of this paper Based on our model described above, we evaluate the power reduction potential for over 10 different approaches applied to an IP-over-WDM backbone network. The baseline of the equipment power efficiency is set at the year 2010. The key novelties of our work with respect to earlier publications are the following.

- First, we provide a formalized methodology to assess the combined power saving potential of various approaches. This assures a large amount of transparency to the power reduction calculations.
- Second, we categorize the power saving approaches by the multiplication factors in our model, and provide a tabulated survey of the most relevant publications for each approach, including the reported power saving potential. This provides the reader with an insight on the consensus of the values reported by different researchers.
- Third, for each approach, we infer a power reduction factor both for a Moderate Effort scenario and a Best Effort scenario. This provides the reader with a plausible range of the saving potential.

Limitations of this paper We should not be mislead by the analytical nature of our methodology to overestimate the accuracy of the final results. As with all models, it has only limited capability in capturing and representing the actual situation it models. Therefore, it makes sense to specifically point out the scope and limitations of this work and approach. First, we focus on the electricity consumption of backbone networks only. We do not consider the embodied energy; the embodied energy is the sum of all energy required to produce equipment. Nor do we consider the opportunities (and issues) associated with using renewable energy sources to power network nodes. The wired and mobile access network are out of scope, as including these would make for an unreasonably large survey; in addition they would require separate models. Data centers are also considered out of scope, although we do briefly touch upon them when discussing caching (Section 5.4.3).

Second, our baseline is an IP-over-WDM network in the year 2010 (with a projection to the year 2020). As such, we do not take into account legacy network equipment and intermediate transport technologies (such as SONET/SDH, or OTN). This is an important limitation; we provide some discussion on the applicability of our findings and potential impact of legacy network equipment in Section 5.5.5.

Third, our breakdown of network power consumption in a number of factors restricts the scope of the power savings approaches we can reasonably capture. For example, we do not specifically identify approaches such as 'multi line rate' versus 'single line rate' transmission, or Elastic Optical Networks (EONs). However, we do think that we capture the majority of approaches reported in the literature.
Fourth, this work remains a survey, and is not an evaluation over a standardized benchmark or baseline. The material we survey is from publications with a wide range of quality control; nonetheless we have tried to use journal papers instead of conference papers when available, and tried to give more weight to results from publications that appeared to be more thorough. We perform a best-effort estimation of the saving potential of various approaches. Evaluating most of the presented techniques on a consistent baseline scenario (topology, traffic matrix, equipment power consumption values) would be the next logic step to do, but is out of scope of this paper.

Organization of this paper For the remainder of this paper, we first describe the earlier mentioned related work in more detail in Section 5.2. We outline our methodology and the associated analytical power model in Section 5.3, before using this model in Section 5.4 to survey a number of power reduction approaches across our baseline. Finally, in Section 5.5 we present the combined power saving potential, discuss its sensitivity to a number of model parameters such as the network hop count, and compare them to the Green Meter results [10]. In Section 5.6 we summarize our findings and complement it with a number of recommendations for the research community, network operators, vendors and policy makers.

5.2 Related work

This section discusses the main works that have been cited in the preceding section. The first paragraph covers general survey papers. The second paragraph covers works that do quantify the isolated and combined power saving potential in backbone networks, and highlights some of their shortcomings. The main works are listed in Table 5.1.

In the work by Restrepo et al. [17], power reduction approaches are broken down into three levels depending on the area of application: (i) on circuit level (such as the use of dynamic voltage or frequency scaling techniques), (ii) on equipment level (e.g., replacing components by their counterpart in the optical domain) and (iii) on network level (network planning for efficiency, e.g. by using optical bypass). In [5], Zhang et al. consider the following techniques to improve the energy efficiency in the backbone network: (i) selectively turning off network elements, (ii) energy-efficient network design, (iii) energy-efficient IP packet forwarding, and (iv) green routing. The power savings for the surveyed works are listed as either 'low', 'medium' or 'high'. In [18], and the earlier survey paper [6], Bolla et al. classify the approaches as either: (i) re-engineering (more energy-efficient network elements, e.g. replacing electronics by optics where possible), (ii) dynamic adaptations (scale power consumption with actual load, e.g. dynamic voltage or frequency scaling), and (iii) sleeping/standby (drive unused network devices to low standby modes). The power savings for the surveyed works are listed throughout the text, but not cross-compared. Bianzino et al. identified in [7] four branches of green networking research, being (i) adaptive link rate (including sleep mode and rate switching), (ii) interface proxying, (iii), energy-aware infrastructure, and (iv) energy-aware applications. The power savings for the surveyed works are not listed. As is clear from this brief survey, many different approaches have been used to categorize power saving approaches. All these categorizations come with their own merits and drawbacks. We feel that categorizing the different approaches according to the factors in a concise analytical model allows for a more insightful estimation to quantitatively assess the potential power savings.

The first work, to our knowledge, that does provide a quantified breakdown of various approaches to reduce the power in telecommunication networks is by Roy [8]; however, the work is rather economical on references. The study by Lange et al. [2] is very holistic and complete in that it does a thorough assessments of the power consumption across the various network segments, expected trends, and opportunities to increase the energy efficiency. However, the quantified saving potential in the backbone is only split up in two high level approaches (technology progress/energy aware systems, and load adaptive network operations), and the associated savings (55% and 20%, respectively) are hard to verify or trace back. The two papers by Kilper [1, 4] span a similar range of holistic topics as the work by Lange, however the saving potential of various approaches is not explicitly quantified. The work by Parker and Walker [9] was the initial inspiration for our study. Next to proposing an absolute energy efficiency metric (dBε) for any ICT system, it presents a notable first effort to provide a synoptic analysis of the performance of 10 different techniques to achieve a 1000-fold reduction in the power consumption of future photonic networks; however, their estimates might be rather optimistic. The Green Meter [10] by the GreenTouch consortium is probably closest in spirit to our work. For backbone networks it identifies and quantifies the power reduction potential of seven approaches, and the combined overall reduction factor. While the fact that some of the approaches are evaluated on a common benchmark presents a rare (but deeply needed) advantage over the other works, this is also stems from the fact that most of the cited works refer to only a few partners in the GreenTouch consortium; as such, the consensus over its findings across the research community outside of GreenTouch cannot be assessed by the reader. The two works by Le Rouzic et al. [11, 12] present an interesting overview of the different backbone saving potentials as identified in the European TREND network of excellence. However, the quoted savings are mainly isolated values; their interplay and the combined saving potential is not given.

Finally, it is worth mentioning the work by Masanet et al. [19]. They estimate the energy use and efficiency potential of U.S. data centers using a methodology very similar to our study. They use a bottom up analytical modelling approach to capture and categorize the total data center electricity demand, and discuss and evaluate a number of efficiency measures. These measures are then presented as reduction factors which can be multiplied to compute the combined electricity saving potential.

5.3 Methodology

Our methodology to assess the impact of efficiency improvements in backbone networks is based on a concise analytical model, which is described by the general form in Eq. (5.1). It expresses the network-wide power consumption associated with an equipment type (such as an IP router) based on a number of factors. Basically, the power rating of the considered equipment (W/Gbps) is multiplied with the amount of traffic in the network (Gbps) and the number of network hops. Three correction factors account for the power consumption associated with general overhead, traffic protection, and overprovisioning. A hop count correction factor is relevant to correctly account for the hops in either the IP layer, or the Wavelength Division Multiplexing (WDM) layer. By considering reduction factors for each of the parameters in Eq. (5.1), we can derive the total reduction potential across an energy-optimized network.

$$P_X = \eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot (H + c_x) \cdot W_x \tag{5.1}$$

where

P_X	total network power consumption for equipment type x (W);
η_{eo}	external overhead factor, e.g. cooling;
η_{pr}	protection factor;

η_{op}	overprovisioning factor;
T	total traffic in the network (Gbps);
H	average hop count in the respective network layer;
c_x	hop count correction factor (0 or 1)
W_x	weighting factor of equipment x, based on its power rating $\frac{P_x}{C_x}$ (W/Gbps) multiplied with an equipment dependent correction factor.

This short overview of our methodology provides us with the required background to outline the (more detailed) remainder of this section.

We first define our baseline, i.e. the reference network scenario over which we will evaluate the different power saving approaches; the equipment power rating values will serve as baseline power weighting factors for the different equipment types (Section 5.3.1). We then discuss the analytical model in more detail and expand on the general form given above (Section 5.3.2). With this established, we detail how this model is used to calculate the impact of efficiency improvements from different contributors on the complete backbone power consumption (Section 5.3.3). Following a short discussion on different ways to express power savings (Section 5.3.4), we introduce two savings scenarios which will provide a lower bound and upper bound for the likely saving potential in backbone networks (Section 5.3.5).

5.3.1 Baseline scenario and reference power rating values

Our baseline architecture is an IP-over-WDM network. As shown in Fig. 5.2, we consider IP routers (which include line cards), long-haul transponders (labelled TXP), Optical Cross-Connects (OXCs), and Optical Line Amplifiers (OLAs). The IP routers switch the IP traffic in the IP layer. The long-haul transponders transmit and receive the optical signal over dedicated wavelengths in fibers. The transponder capacity is 10G, and the fiber multiplexes 40 wavelengths. OXCs provide optical switching capabilities in the WDM layer, by adding and dropping the wavelengths in the different network nodes as required. OLAs are required typically every 80 km and amplify all wavelengths in a fiber. As there is a move from multiple stacked technologies to IP-over-WDM, we leave out other potential intermediate switching technologies, such as Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH), Optical Transport Networking (OTN) and Ethernet. Furthermore, we do not include the power consumption of 3R regenerators (required for optical channels span-

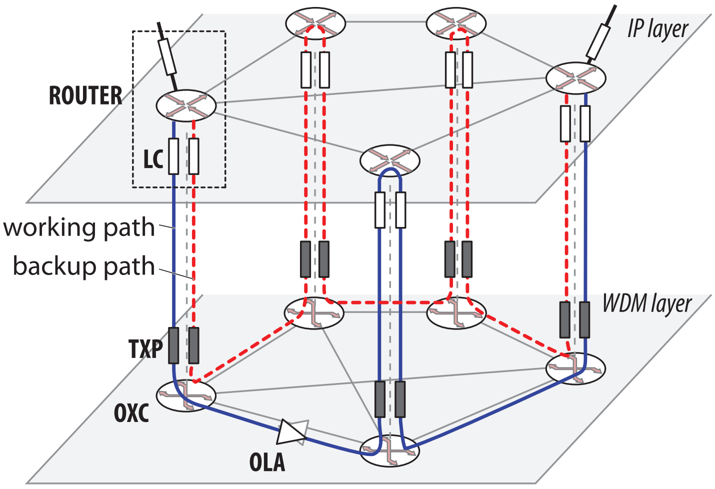

Figure 5.2: *Conceptual model of our generalized IP-over-WDM network architecture. It shows a bidirectional working path (solid lines) and backup path (dashed lines) under a 1+1 protection scheme. (LC = Line Card, TXP = Transponder, OXC = Optical Cross-Connect, OLA = Optical Line Amplifier)*

ning large distances, a typical value is over 1500 km), as their contribution to the total power consumption is marginal [16].

The power rating values we will consider for the four types of equipment are given in Table 5.2. The power rating expresses the power per capacity, in W/Gbps, and is as such an (inverse) measure for the power efficiency of the considered equipment. We will use them as a basis for the weighting factors (also listed in Table 5.2) in the analytical model, as we will explain in the next subsection. The power ratings are distilled from [16] and [20], and are homogenized across three properties. First, they are derived from typical power consumption values (as opposed to vendor rated power, which is typically higher and used for provisioning the power distribution infrastructure). Second, the power rating includes the power associated with the chassis and control overhead (as opposed to only the power for the functional component, which might lead to very low values for e.g. OLAs). Third, the power rating value is that for realistic filling levels of shelves and racks (as opposed to optimal power rating values when assuming maximally filled shelves/racks) [20].

Finally, these values are relevant for the year 2010, which is the reference year in our baseline scenario. The reason we choose the year 2010, is because that is the most recent year for which we have reliable power consumption values available for *all* backbone equipment considered, as provided by our earlier publication [16] which was dedicated to this.

Table 5.2: *Power rating values and weighting factors, reference year 2010 (based on [16, 20])*

Type	Power	Power rating P_x/C_x	Weight W_x
Core IP router (inc. line cards)	-	10.00 W/Gbps	10.00
OXC [(a)]	-	0.46 W/Gbps	0.92
Transponder 10G	50 W	5.00 W/Gbps	10.00
OLA 80 km [(b)]	165 W	0.41 W/Gbps	2.06

[(a)] Power rating 0.46 W/Gbps of an OXC node degree 3, for 40× 10G channels.
[(b)] Power rating 0.41 W/Gbps of a bidirectional OLA for 40× 10G channels.

5.3.2 Analytical power model

The analytical power model that we will use is inspired both by the work from Baliga et al. [3], Kilper et al. [1] and the findings from our earlier work [16].

The total power $P_{BACKBONE}$ in an IP-over-WDM network can be given as the sum of the power consumption in the constituting layers:

$$\begin{aligned} P_{BACKBONE} &= P_{IP} + P_{WDM} && (5.2)\\ &= P_{IP} + (P_{OXC} + P_{TXP} + P_{OLA}), && (5.3) \end{aligned}$$

with P_{OXC}, P_{TXP}, and P_{OLA} respectively being the total network power consumption of the OXCs, the WDM transponders, and the OLAs.

The power consumption of each equipment type is further given as:

$$P_{IP} = \eta_{eo} \cdot \frac{\eta_{pr}}{2} \cdot \eta_{op} \cdot T \cdot (H+1) \cdot \left(\frac{P_{ip}}{C_{ip}} \cdot 2\right) \qquad (5.4)$$

$$P_{OXC} = \eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot H \cdot \left(\frac{P_{oxc}}{C_{oxc}} \cdot 2\right) \qquad (5.5)$$

$$P_{TXP} = \eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot H \cdot \left(\frac{P_{txp}}{C_{txp}} \cdot 2\right) \qquad (5.6)$$

$$P_{OLA} = \frac{\eta_{eo}}{2} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot H \cdot \left(\frac{P_{ola}}{C_{ola}} \cdot \left\lfloor \frac{\text{link length}}{80\,\text{km}} \right\rfloor\right) \qquad (5.7)$$

Note that we use upper case subscripts for the total network power consumption (e.g., P_{IP}), and lower case subscripts to denote a single equipment power consumption value (e.g., P_{ip}).

The external overhead factor η_{eo} accounts for the power consumption due to external cooling and facility overheads in telecom centers, with typically $\eta_{eo} \approx 2$. This value is not applicable to OLAs as they are typically deployed in dedicated outside cabinets without active cooling, so we apply a correction factor $\frac{1}{2}$ to η_{eo} in Eq. (5.7). The protection factor η_{pr} accounts for traffic protection, with $\eta_{pr} \approx 2$ for backbone networks using a 1+1 protection scheme (i.e., all traffic is routed twice on link-disjoint paths). As this is typically done in the WDM layer but not the IP layer (see further), again we apply a correction factor $\frac{1}{2}$ to η_{pr} in Eq. (5.4). The overprovisioning factor η_{op} accounts for the overprovisioning of the network capacity to deal with unexpected traffic spikes and future traffic growth. The traffic factor T gives the total amount of traffic in the network (in Gbps). The hop count H is the average number of hops between processing elements in the respective layer. The reason we have $(H+1)$ for the IP layer is because we also need to account for the client side capacity of the IP router, i.e., towards the access network[3]; for more details see [16]. The power rating factor $\frac{P_x}{C_x}$ expresses the average power per capacity (in W/Gbps) for a given equipment x, as listed in Table 5.2. The factor 2 at the end of the equation accounts for the fact that for each hop the relevant node capacity is required at both the sending and receiving side[4]. The OLA power consumption is a function of the average link length, as an OLA is required every 80 km.

To further generalize on the four equations above, we will combine the factors between brackets and the two $\frac{1}{2}$ correction factors for IP routers and OLAs into a weighting factor W_x, resulting in the following general equation:

$$P_X = \eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot (H + c_x) \cdot W_x \tag{5.8}$$

with $c_x = 1$ for the IP layer, and $c_x = 0$ for the equipment in the WDM layer. Note that is exactly the general form of our analytical model as in Eq. (5.1). The weighting factors W_x are listed in Table 5.2. For the OXCs, and transponders the weighting factor is simply twice the power rating value $\frac{P_x}{C_x}$. For the IP routers the weighting factor is equal to the power rating value. For the OLA contribution in Eq. (5.7), we calculated the weighting factor assuming an average link length of 800 km, which is a reasonable value for backbone networks. In Section 5.5.2 we will look at the impact of the link length on our efficiency improvement estimates, and show that it is almost negligible.

---

[3] In [16] we actually used $(H + 1/\eta_{pr})$ instead of $(H+1)$. However, the slight increase in accuracy does not affect our results in any meaningful way, but does provide additional complexity to calculate the power saving potential. Therefore, we just use the more simple term $(H+1)$.

[4] This approach implies that all traffic is bidirectional, i.e., that there is as much traffic from node A to node B, as from B to A. While this is a simplification, the approximation is adequate enough for our purpose.

5.3.3 Calculating power savings

5.3.3.1 Savings for one equipment type

To estimate the savings for an equipment type x in a backbone network, we can model an improvement in each of the factors in Eq. (5.8) by inserting beta reduction factors (with $\beta \geq 1$).

$$P_{X,impr} = \frac{\eta_{eo}}{\beta_{eo}} \cdot \frac{\eta_{pr}}{\beta_{pr}} \cdot \frac{\eta_{op}}{\beta_{op}} \cdot \frac{T}{\beta_t} \cdot \left(\frac{H}{\beta_h} + c_x\right) \cdot \frac{W_x}{\beta_{pc}} \tag{5.9}$$

Each improvement factor β can be seen as an *approach* that acts independently to reduce the power consumption. The goal of the upcoming Section 5.4 will be to determine feasible β values for a wide number of approaches.

Thus, the total power reduction $\beta_{x,tot}$ for the equipment type x of an improved network consuming $P_{X,impr}$ power compared to the referenced network consuming $P_{X,ref}$ is given by:

$$\beta_{x,tot} = \frac{P_{X,ref}}{P_{X,impr}} \tag{5.10}$$

$$= \frac{\eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot (H + c_x) \cdot W_x}{\frac{\eta_{eo}}{\beta_{eo}} \cdot \frac{\eta_{pr}}{\beta_{pr}} \cdot \frac{\eta_{op}}{\beta_{op}} \cdot \frac{T}{\beta_t} \cdot \left(\frac{H}{\beta_h} + c_x\right) \cdot \frac{W_x}{\beta_{pc}}} \tag{5.11}$$

Note that we can cancel out the factors η_{eo}, η_{pr}, η_{op} and T.

5.3.3.2 Savings across multiple equipment types

The total power consumption in an IP-over-WDM network is the sum of the power in the IP layer, the OXCs, the transponders and the OLAs. To calculate the total power reduction β_{total} across the complete network, we now have:

$$\beta_{total} = \frac{P_{ref}}{P_{impr}} \tag{5.12}$$

$$= \frac{\sum_x^{equip} \left[\eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot (H + c_x) \cdot W_x\right]}{\sum_x^{equip} \left[\frac{\eta_{eo}}{\beta_{eo,x}} \cdot \frac{\eta_{pr}}{\beta_{pr,x}} \cdot \frac{\eta_{op}}{\beta_{op,x}} \cdot \frac{T}{\beta_{t,x}} \cdot \left(\frac{H}{\beta_{h,x}} + c_x\right) \cdot \frac{W_x}{\beta_{pc,x}}\right]} \tag{5.13}$$

Note that the β factors can be different for each equipment type, e.g. β_{pr} can be different for the IP routers and the transponders; therefore we have added an x index to each β factor.

The occurrence of the c_x term is unfortunate, as it makes the result dependent on the hop count H (i.e., we can not factor out all β factors). While for a given topology the hop count will depend on several aspects, such as the routing algorithm and link weights, a good ballpark number of H in a backbone network is 3–4 hops [1, 16, 21]. Therefore we will assume $H = 3$ when calculating the saving potential in the next section. In Section 5.5.3 we will look at the impact of changing the hop count H on our power improvement estimates, and show that it is rather limited.

5.3.4 Expressing power savings

In the above sections, we have modelled power improvements using a *reduction factor* $\beta \geq 1$. In contrast, by far the most common approach used in publications on energy efficiency in networks is to state the *savings percentage* γ (with $\gamma = 1 - \frac{1}{\beta}$). For example, a power reduction with a factor $\beta = 4$ corresponds to a savings percentage $\gamma = 75\%$. It might be interesting to note that another variation is to express the reduction factor in decibel[5] (dB), i.e. on a logarithmic scale as $10 \log_{10} \beta$.

While both the reduction factor β and the savings percentage γ are mathematically interchangeable, the power reductions *intuitively* communicated by them is different. For example, when comparing the power savings $\gamma_1 = 80\%$ and $\gamma_2 = 90\%$, the latter appears to be only a slight improvement over the first one. While this is indeed true compared to the baseline (i.e., the original power consumption), it does somewhat conceal the fact that the 90% savings scenario consumes only half of the power of the 80% savings scenario. Given the historic and projected exponential growth of traffic in the backbone network (see Fig. 5.1), it is important to realize we need and are interested in significant power reductions. A power saving approach (or more likely, any combination of approaches) that can provide 90% savings instead of 80% savings allows for an extra doubling of the traffic while still consuming the same amount of power. In contrast, if we would have identified the power savings using a reduction factor $\beta_1 = 5\times$ and $\beta_2 = 10\times$ respectively, the factor 2 in difference would have been instantly clear.

Therefore, in the upcoming Section 5.4 we will express our estimates of the power reduction potential of the various approaches using β reduction factors (≥ 1). However, when we quote and list existing works (such as in Table 5.6) we will cite the savings as a percentage γ so that the values given can easily be verified in the referenced works. For the reader's convenience, we have listed a set of savings percentages and the corresponding reduction factors in Table 5.3; if required, this can be consulted when going through the subsequent sections.

---

[5] Expressing the power reduction in decibel has the advantage that improvements can easily and intuitively be summed (instead of multiplied), and is for example for this reason used in [9].

Table 5.3: *Conversion table for the reader's convenience*

Savings percentage γ	Reduction factor β
0%	1×
10%	1.11×
20%	1.25×
30%	1.43×
40%	1.67×
50%	2.00×
60%	2.50×
70%	3.33×
80%	5.00×
90%	10.00×

5.3.5 The Moderate Effort and Best Effort scenario

In our analysis in the upcoming Section 5.4, we will distinguish between two scenarios to model the power savings that are possible in backbone networks.

- In the **Moderate Effort scenario** we use small power reduction estimates, corresponding to solutions that are relatively feasible from either a technical or operational point of view. Where the power reduction associated with an approach is unsure, we will consider a likely lower bound that we will model in this scenario.
- In the **Best Effort scenario** we use larger, more aggressive power reduction estimates, corresponding to solutions that would be more challenging to implement from a technical, operational or cost perspective. Where the power reduction associated with an approach is unsure, we will consider a likely upper bound that we will model in this scenario.

Combined, both scenarios provide a range for the achievable power savings potential in backbone networks.

5.4 Approaches to save power

In this section, we discuss several approaches to reduce the power consumption in backbone networks. We do so using the parameters in Eq. (5.1). Fig. 5.3 gives an overview of the different power saving approaches that we will discuss, and how they map to the various factors.

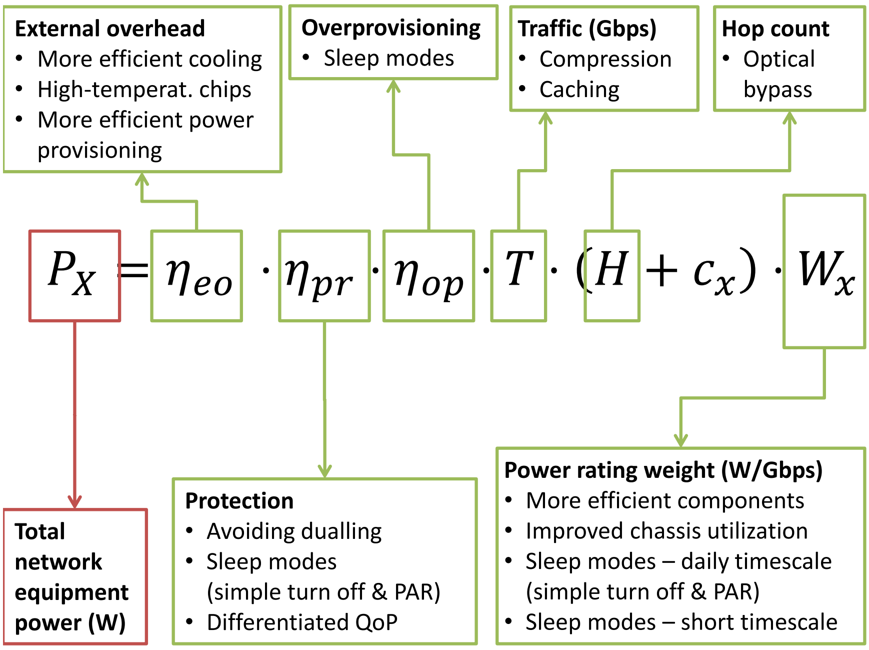

Figure 5.3: Overview of power reduction approaches mapped to the general form of our analytical model

5.4.1 External overhead factor η_{eo}

The external overhead factor η_{eo} accounts for power consumption associated with external cooling and facility overheads in telecom centers. This overhead is commonly characterized by, and also commonly known as, the Power Usage Effectiveness (PUE). The PUE is the ratio of the total amount of power consumed to the useful power consumed[6], and typically has a value of about 2 [24]. In this specific case, this means that for each Watt consumed by useful equipment, such as servers and switches, an additional Watt is consumed through external overhead. In highly optimized and efficiently cooled data centers (much) lower PUE values are possible[7], but this is not yet commonplace. On average, this overhead is made up of two main contributing components [27] [28]: cooling and air conditioning, and efficiency losses in power provisioning (see Fig. 5.4); the contribution of switchgear and lighting is only minor.

We discuss three approaches to reduce the external overhead factor η_{eo}: more efficient cooling systems, high-temperature chips, and more efficient power provisioning.

More efficient cooling A first approach is to increase the efficiency of the premises cooling, i.e. to provide the same degree of premises cooling effect while using less electrical energy. This is a particularly hot topic in data center research, and a wealth of publications on this topic is available. Good overviews are available in [29], [30] and [31]. Examples of such approaches include hot aisle/cold aisle (to avoid mixing both cold and hot air), free cooling (using cold outside air, if ambient temperature and humidity permit), and rack liquid cooling (to improve heat transfer).

High-temperature chips A straightforward alternative to reduce cooling power is to cool less [30]. For example, it has been observed by a notable study from Google [32] that contrary to popular belief, hard disks do not become less reliable when running at higher temperatures. However, this is only possible up to a certain limit. An approach taking the concept

---

[6] While the PUE concept is elegant and simple, in practice there are many intricacies that make it easy to result in different values. To illustrate, while the 2008 white paper that documents the PUE methodology contained only 9 pages [22], the version from 2012 consists of 83 pages [23].

[7] For example, Google states to have reached an annualised average PUE across all their tracked data centers of 1.14 by the end of 2011 (with a minimum value of 1.11 and a maximum value of 1.21) [25]. This value is even more impressive, as they claim that their PUE calculation is more comprehensive than what is done by other players, by accounting also for overhead sources that are typically omitted, such as data center offices and site substation losses. According to the 2012 Data Center Survey by the Uptime Institute [26], the PUE reported by its participants averages between 1.8 and 1.89.

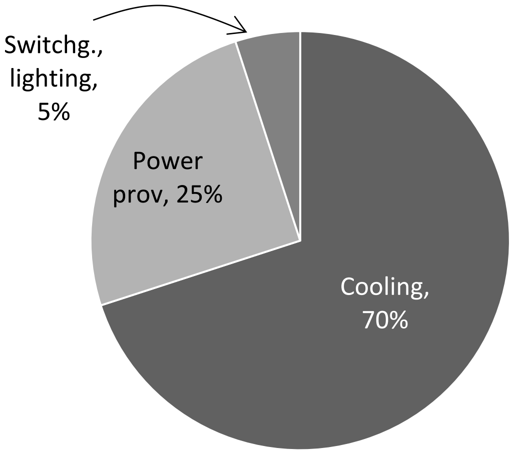

Figure 5.4: *Rough distribution of the external overhead power among its main contributors (source: [8, 27, 28])*

beyond this limit is that of research into the high-temperature operation of integrated circuits [9]. Such an approach further reduces the need for cooling, and as such brings down the external overhead factor η_{eo}.

More efficient power provisioning Power provisioning accounts for roughly 20% to 30% of the external overhead power [8, 28], mainly through inefficiencies in Uninterruptible Power Supply (UPS) units and Power Distribution Units (PDUs) [27]. While the efficiency of a UPS unit can be around 90% at maximum load, its efficiency drops off steeply when lightly loaded—which is very often the case [33]. One reason for lightly loaded provisioning equipment is that such equipment is deployed based on nameplate power ratings (i.e. vendor indication of the maximum power drawn, which is used for dimensioning the power supply systems) of the ICT equipment, which can be substantially higher than the actual peak power [34]. Right-sized provisioning of equipment to the actual peak power of the ICT equipment would reduce the overhead power consumption.

To asses the total overhead power saving potential, we could try to assess the individual potential of each of the above three methods. With cooling comprising typically more than 60% of this overhead ([8, 28]), the biggest gains can probably be achieved there. The study by Roy [8] gives an indication[8]. It estimates that implementing cooling best practices and supplemental high density cooling can save 44% on the required cooling power. Optimizing the power provisioning by replacing legacy rectifiers

---

[8] See Table 6 in [8].

with new generation rectifiers (to increase the peak efficiency), using DC-powered IT equipment, and employing DC ECO mode (to improve rectifier efficiency at lower loads) can save 71% on the power provisioning losses. Combined, this reduces Roy's central office PUE value from 2.14 to 1.63, wich is a reduction of about 1.3× (or 23% savings). However, this value depends on the original PUE value. Therefore, we will go for a more pragmatic approach and lump all approaches together in the PUE factor, and explore how this factor can be improved. To do so, we require (a) a good estimate on the baseline PUE value for telecom network infrastructure in 2010, and (b) a realistic estimate on the improved PUE achievable at these premises.

While there have been a number of studies on the PUE values of data centers (such as [24, 26]), little information is available on a *baseline PUE* for telecom network infrastructure. A first approach would be to just use the average data center PUE value. However, we see a few reasons why the average PUE for data centers might not be representative for those of telecom network infrastructure. First, the PUE of telecom network infrastructure might be worse than those for data centers as network equipment typically has a longer operating lifetime and requires high uptimes, so an overhaul of supporting networking infrastructure to improve the cooling overhead is less easy to occur. On the other hand, the PUE of telecom network infrastructure might be better because it does not suffer from the effect of 'in-house' data centers. In-house data centers are facilities owned and operated by companies whose primary business is *not* computing, and as such efficient cooling is not much of a concern, resulting in relatively bad PUEs. These in-house data centers dominate the data centers in total electricity use and this worsens the average PUE of data centers [35]. In Table 5.4 we list some indicative, public values to determine a reasonable value for telecom network infrastructure PUE in 2010. We have also included some data center values from telecom operators; these values are typically used by operators to highlight their energy-savings efforts, and it seems reasonable to use them as a lower bound for their network infrastructure PUEs. In addition to the limited set of public values in the table, we should note that in private communication with operators, one operator reported an average around 1.5–1.6, while according to another operator values up to 2.0 are not an unreasonable average. Taking all of the above into account, an average baseline PUE value of 1.7 for telecom network infrastructure in 2010 seems a reasonable assumption[9]. Note that this is not too different

---

[9] For our analytical model in Section 5.3.2, we approximated the external overhead factor with $\eta_{eo} = 2$ (instead of 1.7), which we used to apply a $\frac{1}{2}$ correction for OLA equipment in Eq. (5.7). This difference does not affect the end results in Section 5.5 in a meaningful way.

Table 5.4: *Indicative national operator PUE values (including for their DCs) for determining TNI PUE values*

Operator/source	Year	DC PUE	TNI PUE
A new Orange DC [37]	2011	1.3	-
Average Telefonica DC [38]	2013	2.4	-
A new Telefonica DC [38]	2013	1.3	-
Wireline central office [8]	2008	-	2.14
Deutsche Telekom average [39]	2005	-	1.75
Deutsche Telekom average [39]	2009	-	1.53
Deutsche Telekom average [40]	2012	-	1.48
Expert interview national operator X	2013	-	1.5–1.6
Expert interview national operator Y	2013	-	1.8–2.0

from an average data center PUE value of 1.88 for the year 2010 [36].

For determining a feasible *improved PUE* value, even less public data is available. The work by Le Masson [41] applies partly to telecom equipment, and describes an experimental setup of a wall structure to improve the free cooling efficiency. The work reports a 6-month average PUE value of 1.28; however, the overhead power does not include any power provisioning, so the value is rather optimistic. The theoretical lower bound of a PUE is 1.0, and highly optimized data centers do reach values as good as 1.12 in 2012 [25]. However, data centers can be built at locations suitable for very efficient cooling (such as close to a river, or at cold geographical locations), whereas telecom central offices are more constrained by the network topology and have been historically put e.g. in the middle of a city. In addition, the issue of legacy network equipment and the required high levels of availability, as mentioned above, makes upgrades to the supporting network infrastructure more complicated than is the case for data centers. So, it seems unlikely that in the medium term the PUE value of telecom network infrastructure can be better (i.e., lower) than 1.2.

Using both the above derived baseline PUE of 1.7 and the improved PUE of 1.2, we get an improvement factor of 1.4×. We assume a slightly lower Moderate Effort reduction factor of 1.3× (or 23% savings); note that this is in line with the estimation from Roy [8] discussed higher up. For the Best Effort reduction factor we assume 2× (or 50% savings), as the baseline PUE might be higher and/or the improved PUE might be lower than what we derived (with the former situation seeming more likely). Since the PUE is not applicable for OLAs (they are deployed in dedicated outside cabinets without active cooling), the final **Moderate Effort reduction factor** becomes 1.27× (or 21% savings) and the **Best Effort reduction factor** becomes 1.85×

Table 5.5: Reduction factor for external overhead

	More eff. cooling [(b)]	High temp. chips [(b)]	More eff. power prov. [(b)]	(Sub)**total**	
Moderate Effort reduction					
IP	-	-	-	1.30×	
Opt. Swit	-	-	-	1.30×	
Transponder	-	-	-	1.30×	
OLA	-	-	-	1.00×	
Weighted total[(a)]				**1.27**×	(=-21%)
Best Effort reduction					
IP	-	-	-	2.00×	
Opt. Swit	-	-	-	2.00×	
Transponder	-	-	-	2.00×	
OLA	-	-	-	1.00×	
Weighted total[(a)]				**1.85**×	(=-46%)

[(a)] See Section 5.3.3 for calculation details.
[(b)] These columns are empty as we estimated the *overall* PUE reduction instead. See text for more details.

(or 46% savings), as shown in Table 5.5.

5.4.2 Protection factor η_{pr}

The protection factor η_{pr} accounts for the additional power consumption due to traffic protection. Traffic protection is typically employed in backbone networks to achieve high reliability to meet costly Service Level Agreements (SLAs). A typical approach in backbone networks is to provide Shared Path Protection (SPP) protection for the IP layer, and to employ a 1+1 protection scheme for lower layers (also referred to as Dedicated Path Protection (DPP)) [48], whereby for each demand between a source and destination node two link-disjoint WDM connections are set up. Some operators adopt an alternative strategy where, in essence, they deploy the backbone network twice with a certain equipment mix using systems from different vendors, with the goal of having diversification over different vendors. Each of the two networks is dimensioned for full capacity, but

Table 5.6: Protection factor related energy savings reported in publications

Source	Savings	Remarks	Justification
Avoiding dualling			
Parker, 2011 [9]	50%	Probably overly optimistic for all network layers. Viable approach for WDM layer.	-
Sleep modes and Power Aware Routing (PAR)			
Cavdar, 2010 [42]	30–40%	WDM layer only. ILP optimization on COST239 network. SPP. For breakdown over impact of sleep mode and PAR, see Fig. 5.5.	Section V.C and Fig. 4 in [42]. Comparing *EASP* with *SBP* at 250 Gbps (low load) and 750 Gbps (high load).
Jirattigalachote, 2011 [43]	30–40%	WDM layer only. Simulation of dynamic content provisioning on COST239 and USNET. DPP. For breakdown over impact of sleep mode and PAR, see Fig. 5.5.	Fig.3 (a) and (b) in [43]. Comparing *EA-DPP Sleep* with *SP-DPP Total*.
Coiro, 2011 [44]	35%	WDM layer only. Exploits adaptation to daily traffic variations, by turning off links. ILP optimization and heuristic on 18-node random generated network. SPP. Part of the savings are attributed to 'overprovisioning typical of transport networks'.	Fig. 5 in [44].
Bao, 2012 [45]	45–65%	WDM layer only. Heuristic on COST239 and USNET. SPP. For breakdown over impact of sleep mode and PAR, see Fig. 5.5.	Fig 5 and Fig 6 in [45]. Comparing *EASPP* curve (referred to as *PASPP* in text) with 100% baseline.
Musumeci, 2013 [46]	36–45% DPP 29–37% SPP	IP-over-WDM. ILP optimization. Savings reported for NSFNET, but claiming to be similar for COST239. Contribution of PAR to savings unclear. Both DPP and SPP. Upper range of savings for higher traffic volumes.	Table 3 in [46]. Comparing for scaling factor $f = 1$ and $f = 10$ the (a) *DPP sleep-mode* result with *DPP all-ON*, and *SPP sleep-mode* result with *SPP all-ON*.
Differentiated Quality of Protection (QoP)			
Lopez, 2013 [47]	10–20%	WDM layer only. Simulation on Telefónica Spanish core network, for three scenario's with a different QoP class mix. Upper range of savings for higher traffic volumes.	Table 4 in [47], savings (roughly averaged) for traffic 1.56 Tb/s and 23.43 Tb/s.

Acronyms: EASP: Energy-Aware Shared Backup Protection, CPE: Customer Premises Equipment, DPP: Dedicated Path Protection, EA-DPP: Energy Aware Dedicated Path Protection, EASP: Energy-Aware Shared Backup Protection, EASPP: Energy-Aware Shared Path Protection, GPON: Gigabit Passive Optical Network, ILP: Integer Linear Programming, PAR: Power Aware Routing, PASPP: Power-Aware Shared Path Protection, SBP: Shared Backup Protection, SP-DPP: Shortest Path Dedicated Path Protection, SPP: Shared Path Protection.

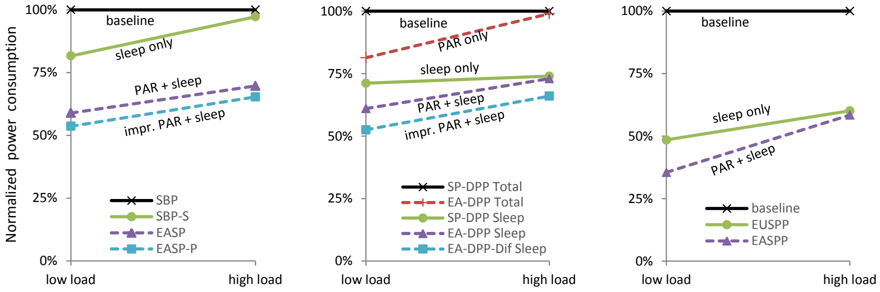

(a) Cavdar **(b)** Jirattigalachote **(c)** Bao

__Figure 5.5:__ Normalized (to baseline) power consumption of several power saving algorithms as reported for the COST239 network by (a) Cavdar [42], (b) Jirattigalachote [43], and (c) Bao [45]. The algorithms have been tagged as either sleep (where port-based equipment is put to sleep), Power Aware Routing (PAR) (where OLAs are put to sleep), or both. The legend shows the algorithm names as they appear in the respective works.

consequently running at 50% utilization, in order to be able to overtake the whole traffic volume. Both of the above approaches result in a baseline protection factor of roughly $\eta_{pr} = 2$, meaning that due to protection the power consumption is doubled compared to a non-protected network.

We consider the following approaches to bring down the protection factor: avoiding dualling, sleep modes, and differentiated QoP. An overview of the main works cited in the subsequent paragraphs is given in Table 5.6.

Avoiding dualling Parker et al. argue in [9] that energy savings of up to 50% of the total network power consumption could be achieved by avoiding the above described network dualling. This would be made possible through increased reliability of network devices, systems and subsytems, and increased software-defined operation of many significant network functionalities. However, this might be an overly optimistic estimation: it is, for example, unclear how more reliable equipment would address the issue of cable cuts. A more viable approach could consist of using more passive forms of protection (passive in the sense of not consuming power). While this would be hard for electrically switched networks (i.e., IP, SDH/SONET and OTN), this is feasible in the optical layer through employing 1x2 passive couplers that duplicate the lightpath from a single transponder (although in that case the service is not protected against transponder failures).

Sleep modes and PAR An alternative approach to reduce the protection factor η_{pr} is to put protection equipment that is serving backup links into a low-power sleep mode. For realistic application, this would require fast sleep and wakeup times. Putting protection equipment to sleep is relatively straightforward for port-based equipment (such as transponders). By using Power Aware Routing (PAR) instead of Shortest Path (SP) routing, additional savings can be achieved by concentrating backup paths and working paths on separate links, in order to be able to put the corresponding OLAs to sleep without being constrained by the presence of working paths. In Fig. 5.5 we show the power saving potential from applying both simple sleep approaches and PAR algorithms as reported by three works that allow to differentiate between these two approaches ([42, 43, 45]). Both Cavdar [42] and Jirattigalachote [43] estimate the combined power saving potential at around 30–40%. The work by Bao [45] reports savings which are more than 20 percentage points higher; however, this work also puts idle working path equipment to sleep[10](which is not to be captured by our protection factor η_{pr}). While the work by Jirattigalachote [43] provides a detailed breakdown across the savings achieved by sleeping only (about 25% reduction, irrespective of the load[11]), and PAR only (about 20% reduction at low load only), these results are not entirely consistent with the work by Cavdar [42] where the savings through sleeping are highly sensitive to the load, see Fig. 5.5.

In contrast to the previous three works that evaluated the power consumption at different loads, Coiro [44] considers a case study where the daily load varies according to a sinusoidal function. He uses only a PAR approach to power off the OLAs, and reports savings of 35%. While the above works only considered the WDM layer, the only work we found that considers both IP and WDM layer is by Musumeci et al. [46]. For the case of DPP, the work reports a 36–45% reduction in energy consumption; there is no breakdown over the impact of sleep mode and PAR. Inconsistent with the previous works, higher savings apply to higher traffic loads.

Summarized, the achievable savings through using sleep mode and PAR for protection equipment are roughly consistent around 30–40%, both for DPP and SPP; however, the savings profile is not always consistent among the various works.

---

[10] Bao [45] is (probably) able to do this because in their baseline scenario the network is not power-optimized at low loads (i.e. the power consumption of the baseline is identical for low and high load), whereas in [42, 43] the baseline power consumption at low load is lower than the baseline power consumption at high load.

[11] The savings are *not* closer to 50% (as one might expect when considering that in a DPP scheme there is a backup path for each working path), because part of the OXC power consumption is considered static and thus not affected by the sleep mode.

Differentiated QoP Ultimately, the customer-demanded level of reliability is a matter of cost. A reduction in protection power consumption could result from having cheaper SLAs that offer (slightly) less reliability with less-demanding customer requirements. This concept has also been branded as differentiated QoP. In [47] Lopez et al. calculate that by using differentiated QoP in a WDM transport network, on average savings of around 10% are possible with respect to DPP, regardless of whether it is a current fix-grid or an envisioned elastic (or flexible grid) network. The savings depend on the traffic load (higher traffic load leads to more savings), and the QoP levels required by the clients.

To assess the total protection factor power saving potential, we again consider a Moderate Effort reduction factor and a Best Effort reduction factor. Even in a conservative estimation, employing sleep mode for protection devices seems to be the most promising solution, as it requires no additional management complexity for the operators. The above reported sleep mode savings of 25% correspond to a reduction factor of 1.33× (applicable to the WDM layer only). Thus, for our Moderate Effort reduction factor we assume that we have a reduction factor of 1.33× for the OXCs, transponders and OLAs. Taking the weight of the OXCs and transponders into account, this results in a **Moderate Effort reduction factor** of 1.14× (or 12% savings), see Table 5.7.

For our Best Effort scenario we assume that (a) that sleep modes and PAR exploit their full potential and can save 40% in the WDM layer (a reduction factor of 1.66), and (b) that differentiated QoP indeed leads to a 10% reduction (a reduction factor of 1.11) at the WDM layer. Note that the combined reduction potential in the WDM layer is 1.84×, which approaches the savings that could be achieved if we could do away with dualling in the WDM layer (or move to passive forms of dualling). As can be seen in Table 5.7 this puts the **Best Effort reduction factor** for protection at 1.29× (or 23% savings).

5.4.3 Amount of traffic T

The yearly growth in backbone traffic T is the major driver behind the continuous increase in power consumption in the backbone network. While the IP traffic growth in the backbone is not as high as it used to be—in the year 2000 traffic was growing at around 125% per year, but has slowed down to around 35% per year in 2010 [59]—it is still projected to grow at a rate that outpaces the improvements in energy efficiency of backbone telecom equipment, as shown in Fig. 5.1.

The traffic estimations in Fig. 5.1 are based on Cisco's two-yearly fore-

__Table 5.7:__ Reduction factor for protection

	Avoiding dualling	Sleep modes	Diff QoP	(Sub)**total**	
Moderate Effort reduction					
IP	1	1	1	1.00×	
Opt. Swit	1	1.33	1	1.33×	
Transponder	1	1.33	1	1.33×	
OLA	1	1.33	1	1.33×	
Weighted total[a]				**1.14**×	(=-12%)
Best Effort reduction					
IP	1	1	1	1.00×	
Opt. Swit	1	1.66	1.11	1.84×	
Transponder	1	1.66	1.11	1.84×	
OLA	1	1.66	1.11	1.84×	
Weighted total[a]				**1.29**×	(=-23%)

[a] See Section 5.3.3 for calculation details.

Table 5.8: Traffic factor related energy savings reported in publications

Source	Savings	Remarks	Justification
Data compression			
Parker, 2011 [9]	50%	Probably overly optimistic	-
Dong, 2012 [49]	45–55%	IP-over-WDM. MILP optimization on NSFNET. Upper range of savings for higher compression efficiencies. Study assumes compression ratios of 20:1 for video and 10:1 for images.	From Conclusion section in [49], which states *‘optimizing data compression [. . .] under the non-bypass approach [saves] up to 45% and 55% of the network power consumption’*
Kilper, 2012 [4]	negative	-	Section VI.B in [4] states that *‘For the case of software-based compression using servers or PCs, uncompressed data transmission ($\sim 10^{-7}$ J/b for ten core hops) was shown to be more efficient than compressed data transmission after including the compression energy ($> 10^{-6}$ J/b depending on the compression ratio [50])’*
Caching			
Lee, 2010 [51] Lee, 2011 [52]	10–60%	Evaluation over 20 content providers, considers complete network (i.e., CPE to core). Lower range of savings for DSL access technology, upper range for GPON. Reason for considerable difference in savings between [51] and [52] is unclear.	From Fig. 4 in [51], *NonCCN* avg estimated at 2750, and *Core100%* avg estimated at 2250, gives 18%. Also states: *‘[Core20% reduces] energy consumption more than 15%.’*. From Fig.4 and Fig 5 in [52] we estimate avg savings of *Core20%* and *Core100%* at 10% and 23% (DSL), and 25% and 60% (GPON) respectively.
Osman, 2011 [53] Osman, 2013 [54]	8–37%	IP-over-WDM. MILP optimization on NSFNET. Upper range of savings for higher demands and improved caching strategy.	Section 5.1 in [54] reports (consistent with [53]) that for the *‘fixed cache size’* strategy daily-averaged savings ranging from 8% (low demands) to 30% (high demands). For the *‘variable cache size’* strategy, the average savings range from 16% to 37%. (Results reported in the Conclusion section are maximum instead of daily-averaged savings.)

Continued on next page . . .

Acronyms: CCN: Content-centric Networking, CPE: Customer Premises Equipment, DSL: Digital Subscriber Line, FT: France Telecom, GPON: Gigabit Passive Optical Network, ISP: Internet Service Provider, MILP: Mixed-Integer Linear Programming.

Traffic factor related energy savings reported in publications (Continued)

Source	Savings	Remarks	Justification
. . . Continued from previous page			
Chiaraviglio, 2011 [55] Chiaraviglio, 2012 [56]	45–65%	MILP optimization for 4 ISP backbone topologies. Reported savings probably capture multiple effects beyond caching, as optimization algorithm also turns on/off servers and network nodes according to daily patterns and exploits server spare capacity, whereas non-energy optimized algorithm does not. Note that, in contrast to [54], low traffic conditions give higher savings, probably because of the above effects.	From Fig. 1.4 in [56], highest savings (low traffic) approx 65%, lowest savings (high traffic) approx 45%.
Modrzejewski, 2013 [57]	10%	Considers complete ISP network (i.e., access to core). Simulation on France Telecom network and Moroccan network.	Table II in [57] states 8.7% and 11.0%. Text also states: '*Energy savings of almost 9% and 11% are possible for the FT and the Moroccan scenarios, respectively*'.
Mandal, 2014 [58]	10–20%	ILP optimization for a hybrid peer-to-peer CDN, considers complete network (i.e., end-user, over access to core). Simulation on a US-wide IP backbone network and passive optical access network. Upper range for more popular content schemes.	Conclusion in [58] states that '*In some cases, our schemes can moderately reduce both server load as well as energy consumption (10%–20%)*' However, the conclusion also reports an increase of energy consumption in other cases.

Acronyms: CCN: Content-centric Networking, CDN: Content Distribution Network, CPE: Customer Premises Equipment, DSL: Digital Subscriber Line, FT: France Telecom, GPON: Gigabit Passive Optical Network, ISP: Internet Service Provider, MILP: Mixed-Integer Linear Programming.

casts [59–62]. The values for the year 2010 and 2011 are taken from [59], while the most recent publication ([62], May 2013) provides a forecast for 2012 to 2017, and estimates the CAGR for the same period at 23%. We have extended this growth rate to the year 2020 (dotted line). As the annual projected growth rate is slowing down (at a rate of about 4 percentage points each year, based on [59–62]), we have also plotted such a more likely, lower growth at 16% per year, from 2014 to 2020 (dotted line, labelled 'adjusted').

As can be seen in Fig. 5.1, even for the low traffic projection the backbone traffic increases 8-fold for the period 2010 to 2020. What are potential approaches to bring down the amount of traffic in the backbone network? The straightforward approach would be to set (reduced) quotas on the amount of traffic granted to each customer. However, this seems very unlikely from an economic perspective. Approaching the issue from a technical perspective, we will look at data compression and caching. An overview of the main works cited in the subsequent paragraphs is given in Table 5.8.

Data compression Data compression, or source coding[12], encodes information in such a way that it requires fewer bits than in the original representation. Parker et al. [9] estimate its potential for power reduction in photonic networking to be up to 50%. However, there are three pitfalls to be aware of. First, it seems unlikely that multimedia content (the bulk of global IP traffic [1]) can be further compressed, in or at the edge of the network. Multimedia content is already heavily compressed, for example audio mp3 compression, and video H.264 compression (used by default by, amongst others, YouTube and High-definition Television (HDTV) broadcasts). In [49], Dong et al. estimate the power saving potential through content compression at about 50%, however, they assume that video and images can be compressed 20 and 10 times respectively, which seems unlikely given that both media formats already widely employ compression at the application side. Second, compression and decompression at transmit and receiving side comes at a processing, and thus energy, cost. Kilper even states in [4] that uncompressed data transmission through the core is an order of magnitude more efficient than compressed data transmission.

---

[12] Source coding should not be confused with *channel coding* and *network coding*. Channel coding in effect *adds* redundancy (instead of removing redundancy, as source coding does), to reduce the bit error rate in noisy communication channels. Network coding is a technique to attain the maximum possible information flow by algorithmically combining packets for transmission. There are studies available that look into network coding to improve the energy efficiency in wireless (ad-hoc) networks [63] and wired PONs [64], but almost none for backbone or transport networks. A single non-peer reviewed study [65] indicates energy savings in the order of 20% for two backbone networks, but states 0% savings for full mesh topologies. Given the limited amount of information available, we do not consider network coding at this point.

That said, the influence of the compression energy could drop with continuously more efficient Digital Signal Processings (DSPs) units [9, 66], and become negligible for content that is accessed by many users over a period of time. Third, any advances in compression techniques will probably be cancelled out by encoding more information and new modalities (such as stereoscopic view, or a higher dynamic range for audio or video) into bit streams (i.e. Jevons paradox). To illustrate this, the historical evolution of video compression factors shows that for the latest main video compression standards, roughly each of them compressed twice as much [67]; the newest upcoming standard (HEVC), is again expected to continue this trend. However, despite this, the average bitrate per video stream has *not* consistently decreased, with a move from Standard-definition Television (SDTV) to HDTV (720 pixels), and then to HDTV (1080 pixels) [68][13].

Caching Another technique to decrease the amount of traffic in backbone networks is the use of caching. The increase of media-rich internet content has lead to high bandwidth requirements for content served to multiple destinations. While the technique of caching is already well established—both at the client side, as well as between client and servers through intermediate proxy servers—Content Distribution Networks (CDNs) are the next logical step. A CDN is a large distributed set of servers deployed throughout the network, with content from the place of origin replicated to the other servers [69]. The main goal of CDNs is to increase availability and performance by serving requests from a server closer to where the request originates, but telecommunication service providers also deploy them to reduce the demand on their backbone. While caching content obviously consumes extra power, for the case study in [54] it is estimated that with optimal cache sizes a reduction of above 30% of the total power consumption (IP-over-WDM network and caching) can be achieved. The important issue to note here is that these savings are expressed taking into account both the reduced energy consumption in the network and the increased energy consumption in the caches. For the sake of argument we attribute all the savings to the network. Modrzejewski performs a similar study in [57]. For two realistic networks of national Internet Service Providers (ISPs) their algorithm to optimally decide where to cache content inside the ISP network predicts about 10% energy savings across the complete network, i.e. from access to core. The work by Lee et al. [52], that has coined the term 'content centric networking', considers network nodes that act as content

---

[13] Ultra high definition resolutions beyond HDTV are emerging, notably 4K UHD (3840x2160 pixels) and 8K UHD (7680x4320 pixels), but we did not find consistent and comparable data on the evolution of the average bitstream per video stream.

caches themselves. They calculate potential savings over a network stretching from the customer side to a Tier 1 ISP core network. Evaluating for the top 20 content providers, they find that 10% to 60% can be saved over the complete network. However, as the large variation is very sensitive to the access network technology (DSL/GPON), this hints that the energy savings in the core will probably be rather limited. A very high saving potential is reported by Chiaraviglio et al. [56] (see Table 5.8), but the evaluation captures savings effects from non-caching approaches as well, so these results are hard to interpret with respect to pure caching. Finally, the use of peer-to-peer caching could potentially reduce the caching power consumption overhead (and thus increase the overall savings), as such a scenario would exploit caching in end-user devices that are already on anyway, instead of integrating entirely new caches into the system as above. An important issue in such a pure peer-to-peer caching scenario is that the energy consumption should not be migrated solely to the end-users, which would be very attractive for network operators but not necessarily reduce the overall (system-wide) energy consumption, as pointed out by Feldmann et al. [70]. A recent work exploring a hybrid solution, where a common CDN is combined with a peer-to-peer caching solution to obtain a more optimal overall system, is by Mandal et al. [58]. The system-wide energy savings (network+caches) are reported as 10–20% in some cases; however other cases are reported as requiring an equal amount or even more energy. Furthermore, a breakdown of the power consumption savings and increase over the different network sections is not given.

Following the above findings, practical approaches to reduce the amount of traffic in the core network seem to be limited. Therefore, for the **Moderate Effort reduction factor** we assume no savings, i.e. a reduction factor of 1.0×. For our Best Effort scenario, we assume that indeed a reduction of about 10% of energy consumption can be achieved both through compression and energy-optimal content caching. This means that we have a **Best Effort reduction factor** of 1.23× (or 19% savings). Both values are summarized in Table 5.9.

Table 5.9: *Reduction factor for amount of traffic*

	Compression	Caching	**Total**	
Moderate Effort reduction	1	1	**1.00**×	(=-0%)
Best Effort reduction	1.11	1.11	**1.23**×	(=-19%)

5.4.4 Power rating P/C (part I)

Table 5.10: *Power rating factor related energy savings reported in publications (part I)*

Source	Savings	Remarks	Justification
More efficient components (annual energy-per-bit reductions)			
Neilson, 2006 [71]	20% p.a.	Trend up to 2006 for high-capacity routers. Note that our trend line in Fig. 5.7 gives 29% p.a., whereas Neilson writes 20% p.a.	Caption of Fig. 4 in [71].
Han, 2010 [72]	15% p.a.	Trend from 1992 to 2008 across 3 generations of Fujitsu optical transport platforms (ADMs → MSPPs → POTPs)	Section II in [72]: *'power consumption for each transported gigabit per second has decreased over the past 16 years, from 84 W/Gb/s for an ADM to 6.7 W/Gb/s for a POTP'*
Tamm, 2010 [48]	13% p.a.	Trend and projection for 2005 to 2020. Applies to routers, packet switches, SDH-XCs, and OTN-XCs.	Fig. 7 in [48].
Lange, 2011 [2]	14% p.a.	Trend and projection for 2002 to 2020 based on routers from 3 different vendors.	Fig. 10 in [2].
Improved chassis utilization			
Ceuppens, 2009 [73]	55%	Difference in power rating of specific Juniper router chassis filled with 1 slot compared to all slots filled.	Slide 6 in [73] for MX960 Juniper router: ECR with 1 line card is 18 W/Gbps, with 11 line cards is 8 W/Gbps.
Van Heddeghem, 2012 [20]	46% (IP eq) 46% (OLA) 9% (WDM)	Savings from suboptimal chassis/shelf filling over usage lifetime for IP equipment, OLAs and WDM terminals, respectively.	Fig. 6 in [20].

Acronyms: ADM: Add/Drop Multiplexer, ECR: Energy Consumption Rating, MSPP: Multiservice Provisioning Platform, OTN: Optical Transport Networking, POTP: Packet Optical Transport Platform, SDH: Synchronous Digital Hierarchy, XC: Cross-Connect.

The power rating factor P/C expresses equipment power consumption as the power per unit capacity, e.g. 5 W/Gbps. While we defined the denominator as the equipment capacity, we will relax this constraint and also include in this section any approaches that improve the power consumption under varying load. The baseline values that we assumed for the various components are listed in Table 5.2.

Fig. 5.6a shows how the power consumption of network equipment

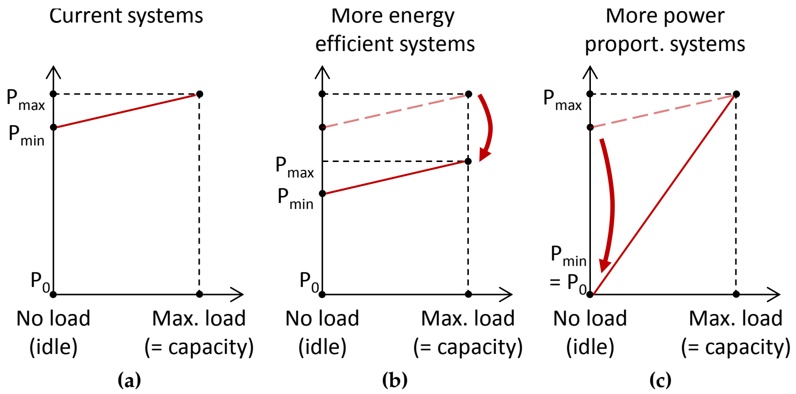

(a) **(b)** **(c)**

Figure 5.6: *Power scaling with load and capacity. (a) For typical network equipment the power consumption when idle is still close to the that at maximum load. (b) More energy-efficient equipment consumes less energy at maximum load. (c) Perfectly power proportional equipment scales linearly with the load.*

typically scales with varying load. Note that the maximum load is equal to the capacity of the equipment. Several works have identified that for network switching equipment the power consumption when idle is around 90% of that at maximum load. Chabarek et al. [74] measured the power consumption of two Cisco routers in different configurations and under various loads. They observed that while the configuration (chassis and line cards used) significantly influences the power consumption of the router, the load has limited impact. These observations are confirmed by Mahadevan et al. [75], and in our earlier work [20]. In addition, in [20] we showed that this is also the case for OLAs. While to our knowledge there is no public data available for other optical equipment, such as WDM terminals, transponders, or OXCs, it seems to be accepted as a fact that this trend holds there as well.

There are two general categories of approaches to reduce the power consumption shown in Fig. 5.6a, and thus the power rating factor.

- The first category, shown in Fig. 5.6b, focuses on reducing the power associated with the maximum load, i.e., the device's capacity. In this category, we will consider the *inherent component energy efficiency improvements* that are observed for communication equipment year after year, and *improved chassis utilization* through better chassis and shelf filling.

- The second category, shown in Fig. 5.6c, focuses on making the equipment more power proportional, i.e. the power scales better with the actual load. The typical techniques to do so involve putting (sub)components or systems to *sleep*. This can be applied on *short time scales* (in the order of packet-level transmission) and on *long time scales* (in the order of hours and days).

For ease of reading, we will discuss the first category (more efficient components, and improved chassis utilization) in this section. The second category (sleep mode approaches) will be moved to, and discussed in, the subsequent section (Section 5.4.5). An overview of the main works cited in this section is given in Table 5.10.

More efficient components Telecommunication equipment becomes more power efficient each year, largely driven by more energy-efficient CMOS technology. However, as already observed by Kilper [1] and Tucker [76] the rate of improvement has been slowing down. While up to 2006 the energy-per-bit reduction (i.e., J/bit or W/Gbps) in routers was around 20% per year [71], it has slowed down to around 13% per year more recently [48], see also Fig. 5.7. In [1], it has been argued that this efficiency improvement might further slow down as a result of practical limitations inherent to CMOS transistor design. Koomey et al. [66], on the other hand, argue that the efficiency improvement trend for electronic processing equipment will continue through clean slate design, similar to what was observed a few decades ago with the transition from vacuum tubes to discrete transistors and subsequently to microprocessors. Note that the reported efficiency trends in Table 5.10 and Fig. 5.7 apply mainly to (opto)electronic equipment. Components that are dominantly optical (such as OXCs and OLAs), improve at a much slower rate [1], although we did not find credible values for these two components. Given the above observations, a 13% per year reduction in energy-per-bit for electronic and optoelectronic backbone equipment (IP routers and transponders) in the time frame 2010 to 2020 seems reasonable. For dominantly optical components (OXCs and OLAs) we assume no yearly improvements. Note that a 13% per year reduction in energy-per bit (i.e., J/bit or W/Gbps) corresponds to a 15% improvement in energy efficiency (i.e., bit/J or Gbps/W), since $1/(1-0.13) = 1.15$. These are the values reported in Table 5.12.
For our complete backbone communication stack, the resulting yearly reduction in energy-per-bit (W/Gbps) then becomes 11%, or a corresponding 13% efficiency improvement (Gbps/W) which is the improvement rate plotted in Fig. 5.1. It is important to point out that this is an idealized improvement rate, and would only be achieved in practice if all equipment in the network is continuously replaced by the latest (i.e., most efficient)

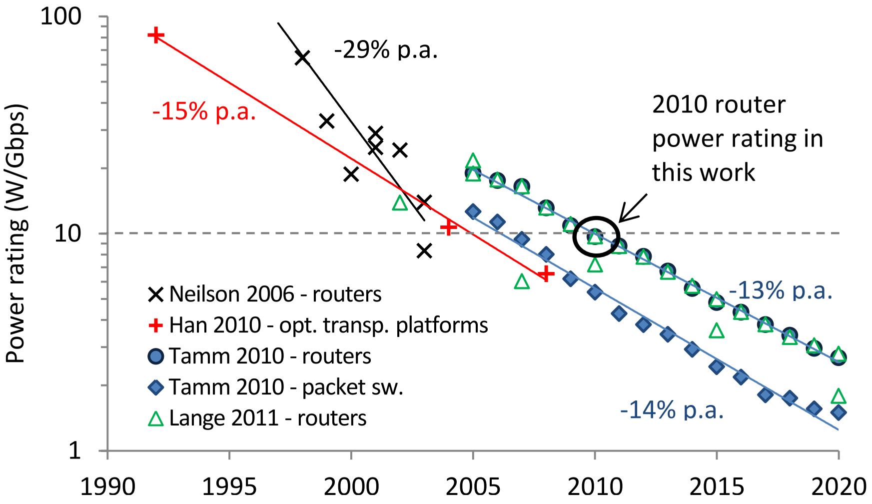

Figure 5.7: *Evolution of power rating values of various telecom equipment as given by Neilson [71], Han [72], Tamm [48], and Lange [2]. Exponential trend lines have been added. Note that some data points are actual values, while others are projections; the publication year might give an indication.*

generations, which is obviously unrealistic[14].

At this point, it should be noted that the yearly energy efficiency improvement of network equipment is the only power saving approach in this study that is time-dependent. All other approaches, both those already discussed and those yet to be discussed, are *static* (once-only) approaches; their full energy reduction potential can only be applied once, and not reused. Because of the time-dependency of more efficient components, we will treat the impact of a 11% p.a. reduction in energy-per-bit (= 13% efficiency improvement) separately in Section 5.5.1.2.

Improved chassis utilization While Chabarek et al. already reported in [74] that 'from a power-aware perspective, it is best to [. . .] maximize the number of line cards per chassis.', to our knowledge the work by Ceuppens [73] is one of the first public works that considers the impact of slot filling levels on the *power rating* value of an IP router. Based on measurements on an MX960 Juniper router, he finds that 'chassis utilization below 30% significantly affects [the power rating]', i.e. the chassis with only one line card is power rated at a value twice as high as the one filled with all line

---

[14] In this context, the work by Parker et al. [77] provides an interesting study on energy-efficient upgrade paths for new router generations through mast-slave configurations; however, the work applies to edge routers instead of backbone routers.

cards. In [20], we have considered this issue more in detail, and looked into the impact of equipment filling levels on the power rating value of both IP equipment and WDM equipment. Equipment deployed in the field is not always optimally filled; but instead often starts off with an almost empty chassis which, over time, is filled with more and more line and control cards. As a result, power rating values will only approach their optimal (lowest) value towards the end of the equipment's life, when the chassis overhead is shared by the maximum number of functional components. The results in [20] indicate that the optimal power rating—achieved when the rack is at maximum capacity, i.e., the power rating value typically assumed in power models—in some cases needs to be corrected by close to a factor 2 (i.e, twice as worse). Particularly, the study found that IP routers over their lifetime are a factor of 1.85 more inefficient than at optimal (i.e., fully filled) capacity. Otherwise said, if the power rating of IP routers would scale better with the filling level, they would be 1.85 times more efficient (or save 46%) than is currently the case. As this value is sensitive to the lifetime assumptions, we assume a slightly lower **Best Effort reduction factor** at the IP layer of 1.5× (or 33% savings). Based on the same work [20], we assume for OLA equipment a reduction factor of 1.5× (or 33% savings), and a reduction factor of 1.1× (or 10% savings) for WDM terminals[15]. Note that the WDM terminals in [20] include the transponder power consumption, so we apply the same factor there. For the **Moderate Effort reduction factors** for the above three layers, we assume no savings, i.e. a reduction factor of 1.0×. For a summary, see Table 5.12 in the next section.

5.4.5 Power rating P/C (part II)

This section discusses the sleep mode approaches (on a daily time scale, and on short time scale) to improve the power rating factor P/C. It is a continuation of the previous section, but has split up for ease of reading. An overview of the main works cited in this section is given in Table 5.11.

Sleep modes on a daily time scale A popular research topic to improve energy efficiency is the usage of sleep modes, and examples of such works are ample (see the surveys [6, 7]). We already discussed it in the context of protection as well (Section 5.4.2). The general idea is based on the fact that even in core networks communication equipment is not always working at maximum load. However, power consumption of communication

---

[15] In contrast with the IP equipment, for OLAs and WDM terminals we don't downward adjust the savings reported by [20] in Table 5.10, as they are not so sensitive to a lifetime parameter as is the case with the IP equipment.

Table 5.11: *Power rating factor related energy savings reported in publications (part II)*

Source	Savings	Remarks	Justification
Sleep modes – daily time scale			
General			
Lange, 2010 [78]	19–38%	Models theoretical upper bound of savings based on three daily traffic profiles (layer agnostic). Higher savings apply to perfect power proportional networks, lower savings for more realistic stepwise adaptivity.	Section 4 in [78], savings for aggregated (sinusoidal) traffic.
Simple turn off			
Liu, 2011 [79]	86% (65%)(a)	IP router ports only. Savings for turning off sublinks at low utilizations in a daily traffic profile, on synthetic Internet2-based network. Reported savings also capture (probably large) effect of exploiting overprovisioning.	Section VII in [79].
Idzikowski, 2011 [80], (Idzikowski, 2013 [81])	2–35% (2–26%)(a)	IP line cards only. MILP optimization for three backbone topologies. Savings for turning off sublinks of link bundles ('parallel line cards') at low utilizations. Upper range of savings for higher demands and more gravitational traffic models.	In [80], comparing estimated daily average of *'FUFL'* with *'Static Base Network'* in Figures 6(a) and (c), and Figures 9(a) through (d). Table 4 in [81] reports 15% for MUELL=0.5.
Power Aware Routing			
Chabarek, 2008 [74]	2–65%	Savings *probably* for router chassis + line card. Savings for rerouting/grooming at low utilizations in a daily traffic profile. High savings for single-port line cards and dense networks.	Table IV in [74].
Restrepo, 2009 [17]	10%	IP layer only. MILP optimization in 50-node core network for EPAR based on five distinct energy profiles. Exploits traffic differences among different nodes.	Fig 6 in [17], for *On-Off* and *Log100* (approaches profile of current routers)

Continued on next page . . .

Acronyms: EPAR: Energy Profile Aware Routing, FUFL: Fixed Upper Fixed Lower, MILP: Mixed-Integer Linear Programming, MUELL: Maximum Utilization of Each Logical Link, PAR: Power Aware Routing.

(a) Estimated savings at IP layer, using the finding from [82] that line cards represent about 75% of the IP power consumption.

Power rating factor related energy savings reported in publications (part II – Continued)

... Continued from previous page

Source	Savings	Remarks	Justification
Zhang, 2010 [83]	30%	IP layer only. MILP optimization on NSFNET case study. Savings for turning off chassis and line cards at low utilizations in a daily traffic profile.	Section 3 in [83], daily averaged power savings: *'The power savings varies from 0% to 68.7%, with an [daily] average value of 29.8%'*. Study considers router chassis and line cards.
Fisher, 2010 [84]	40–80% (30–60%)(a)	Probably only applicable to IP router line cards (power model not given). ILP optimization and heuristics. Savings for turning off sublinks in link bundles at low utilizations in a daily traffic profile. High savings apply to larger bundle sizes. Reported savings also capture (probably large) effect of exploiting overprovisioning.	From Fig 2 (Abilene) and 3 (two synthetic topologies) in [84].
Idzikowski, 2013 [81]	39–45% (29–34%)(a)	IP line cards only. Survey and benchmark of 6 different PAR approaches (covering the works [85–89]) exploiting daily traffic variations on France Telecom reference scenario.	Table 4 in [81] for MUELL=0.5. Higher values of MUELL seem undesirable for network operators.
Sleep modes – short time scale			
Nedevschi, 2008 [90]	40–60% (30–45%)(a)	IP line cards only. Savings for inter-packet sleeping at low utilizations using a buffer-and-burst approach. Packet-level simulation on two realistic topologies and traffic workloads (Abilene and Intel), assuming at packet-level a mix of Pareto-flows and constant bit-rate traffic . Upper range of savings for low utilizations.	Figures 14(b) and (d) in [90] are representative for current equipment with high idle power. Savings from 60% (utilization=10%) to 40% (util.=30%)
Reviriego, 2010 [91] Reviriego, 2011 [92]	-	Evaluates the feasibility of Energy-Efficient Ethernet (EEE) on optical high-speed links.	-

Acronyms: EPAR: Energy Profile Aware Routing, FUFL: Fixed Upper Fixed Lower, MILP: Mixed-Integer Linear Programming, MUELL: Maximum Utilization of Each Logical Link, PAR: Power Aware Routing.

(a) Estimated savings at IP layer, using the finding from [82] that line cards represent about 75% of the IP power consumption.

equipment remains almost independent of the actual load ([20, 74, 75]), as illustrated in Fig. 5.6a. As such, making the equipment or system more power proportional by shutting down (sub)components when it is not in use could lead to substantial overall savings. Before expanding on the two different sleep mode approaches to improve power proportionality, it is instructive to look at the daily traffic variation. As shown by Lange et al. in [78], the traffic volume in communication networks varies considerably over time. For aggregated traffic (as opposed to service-specific traffic), the variations range from peak values typically in the evening, to off-peak values as low as 25% in the morning (see the inset in Fig. 5.8). Other actual traffic matrices showing this daily variation can be found in [81] (France Telecom network), [93] (anonymized), and [94] (showing the interesting fact that European network traffic reaches lower off-peak values compared to U.S. traffic). The aggregated traffic in the core can be approximated by a sinusoidal curve, and for current networks the minimum traffic value is in the range of 25% of the peak traffic volume [78]. Note that from a theoretical perspective, if the load exhibits a perfectly sinusoidal variation over the day, with its minimum being zero during off-peak times, the savings will be upper bounded to 50%. More generally, if the off-peak load is α instead of zero, then the theoretical savings are $S(\alpha) = 0.5 \times (1 - \alpha)$, see Fig. 5.8. With the off-peak load being estimated in current core networks at α=25%, the theoretical upper bound to the savings are thus 38%. For more realistic stepwise adaptivity to the load, the savings are up to 19% [78]. Similar to what we saw for exploiting sleep modes for protection equipment (Section 5.4.2), we consider a relatively straightforward *simple turn off* strategy, and a more elaborate *PAR* strategy to maximize the equipment that can be put to sleep; both discussed next.

The most simple approach to exploit this theoretical potential is to *simple turn off* equipment when it is idle. Visualized on Fig. 5.6a this would mean that the power of the relevant equipment at idle drops from P_{min} to P_0. If this can be done stepwise for a number of subcomponents (such as the line cards in an IP router), then we can gradually approach the power proportionality shown in Fig. 5.6c. The premise for being able to simply turn off links, is that one logical (IP) link is actually a 'bundled link'[16] that consists of multiple aggregated sublinks. For example, a 40 Gbps bundled link may be realized through four 10 Gbps interfaces or sublinks. The study by Fisher et al. [84] claims that 'in today's backbone networks, a vast majority of links would be bundled, with bundles consisting of two to approximately

---

[16] Link bundling is also referred to under various other umbrella terms such as link aggregation, link bonding, link teaming and port trunking. The IEEE 802.1AX-2008 standard uses the term 'link aggregation'.

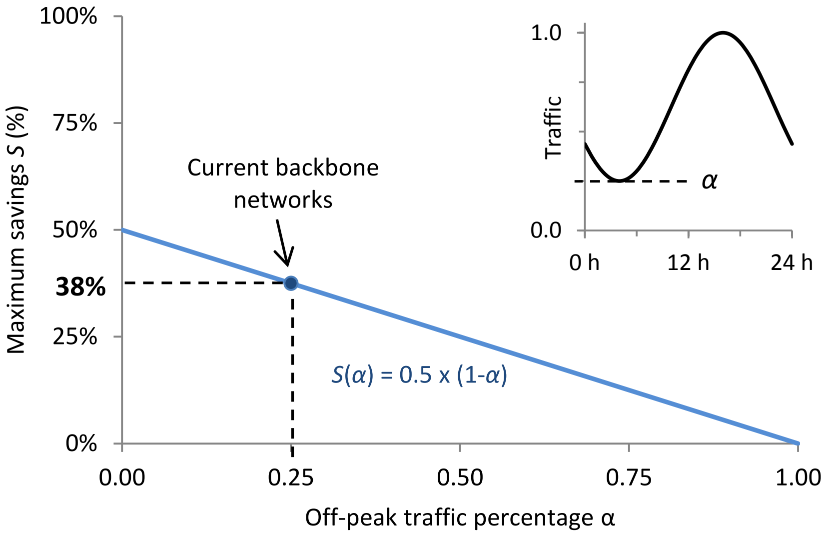

Figure 5.8: *Theoretical achievable sleep mode savings S for a sinusoidal traffic variation over the day. The current off-peak traffic is estimated by [78] at* $\alpha = 0.25$*, which corresponds to maximum theoretical savings of 38%. The sine curve has been phase shifted to approximate the traffic curve from [81] with a minimum around 5am.*

twenty cables [i.e., sublinks], a majority between the two extremes'. It is argued in the same work that the drivers behind link bundles are both (a) capacity requirements exceeding the fastest available link technology, and (b) the fact that capacity upgrades are often realized by adding new links alongside existing ones (rather than replacing the existing equipment with a higher-capacity link). Doverspike et al. [95] mention a third driver for link bundling, which is resilience and consequently network stability; if one of the component links fails the bundled link remains up and a failure-driven topology update is not required. Unfortunately, Fisher et al. [84] do not provide actual data (such as link bundle counts for real operators) to ground their—otherwise plausible—claims. Two works that consider a simple turn off strategy of sublinks (without changing the logical topology) are by Liu and Ramamurthy [79] and Idzikowski et al. [80]. Liu and Ramamurthy [79] report 86% IP port savings by putting sublinks to sleep following daily traffic variations. As line cards represent about 75% of the IP router power consumption [82], this would mean that in the IP layer savings around 65% are possible. The reason for this very high savings percentage (much higher than our indicative, theoretical upper value of 38%, as derived earlier) is that they exploit overprovisioning in the network; they report for a specific link that the 'utilization is commonly less than 1%’. We should be careful not to assign the potential saving effects of exploiting overprovisioning, as this is often installed with good reason (see further). The study by Idzikowski et al. [80] uses a similar technique in its ‘FUFL’ algorithm to turn off sublinks (referred to as ‘parallel line cards’) at low utilizations, resulting in savings ranging from 2% to 26% (IP layer savings) depending on the demand size and gravity of the demand model. This study *does* consider keeping the overprovisioning factor (captured by the Maximum Utilization of Each Logical Link (MUELL) parameter) identical in both the baseline scenario as well as the sleep scenario. So these estimations are more useful for our purpose. The practical advantage of the simple turn off strategy lays in the fact that all decisions can be taken locally and require no logical topology changes, which is especially attractive for network operators where network stability is critical. A significant drawback includes the fact that the actual saving potential relies heavily on the amount of bundled links being prevalent in current and future core networks[17].

A more elaborate approach, which we will refer to as *Power Aware Routing (PAR)*[18], is to not wait until a line card is completely idle, but instead pro-actively reroute traffic so that traffic of lightly loaded links is moved to links with spare capacity. In practice this means that traffic is groomed. This can be done periodically (e.g. multiple times per day, or even per hour) in response to the demands changing throughout the day, resulting in a reconfiguration of the logical topology. The obvious gain of the extra effort involved in PAR is to increase the amount of energy savings that are possible in response to these traffic variations. An important aspect here is the link overprovisioning, also called the maximum link load. Network operators overprovision the capacity of backbone links by a factor of 2 or even more ([1, 97]) to account for peak-to mean traffic variations, unexpected traffic spikes (e.g. in response to a major news event), and future traffic growth (see Section 5.4.7). Any study evaluating the energy saving potential of PAR-like approaches should make sure that the overprovisioning factor is kept identical in both the baseline scenario and the power optimized PAR scenario to allow for a fair comparison. As was shown in Figure 17 of [86] and in Table 4 of [81] (with the overprovisioning factor labelled MUELL), increasing the overprovisioning factor only in the optimized scenario results in significantly higher energy savings (more than 20 percentage points), with the hidden cost being a reduction in Quality of

---

[17] It is interesting to point out that we identified the same high sensitivity to the ratio of the required link capacity over the available interface capacity in [96] (referred to as demand/linerate ratio) in the context of optical bypass.

[18] Note that similar approaches have been referred to in other works as Energy Aware Routing (EAR), Energy Aware Adaptive Routing (EAAR) or Green Routing.

Service (QoS). The issue in some studies is that this influence is not always properly accounted for or clearly indicated[19], which might lead a reader to interpret results overly optimistic. In the context of sleep mode approaches to save energy (both simple turn off, and PAR), the work by Idzikowski et al. [80] is specifically interesting as it both (a) has a good survey on earlier related work highlighting some important open issues in these works, and (b) investigates in a consistent manner the sleep mode saving potential taking into account consistent overprovisioning. Furthermore, some of the algorithms proposed in [80] are evaluated on a testbed by Tego et al. [98], which so far has only rarely been done for research in this area. For the saving potential of PAR approaches we will however refer to the more recent work by the same author [81], as it evaluates different PAR algorithms (notably from [85–89]) across a consistent scenario (topology and traffic matrix). The reported savings for the different algorithms are in a very narrow range (29% to 34%, IP layer savings, see Table 5.11). The other PAR-related works we have listed in Table 5.11 ([17, 74, 83]) report savings consistent with that finding. A notable exception is Fisher et al. [84] who report much higher savings (40% to 80%, *probably* applicable to IP router line cards, which would correspond to IP layer savings around 30% to 60%). This is most likely because the work also exploits overprovisioning.

Summing up, the savings attributable to sleep modes that exploit daily traffic variations, the 'simple turn off' strategy seems attractive both from the real savings it can offer and the limited added technical and operational complexity required. Thus, for our **Moderate Effort reduction factor** at the IP layer we assume 1.15× (or 13% savings), as an average of the IP layer savings estimated by [80]. With the more operationally challenging PAR we assume for the **Best Effort reduction factor** the upper bound from [81], i.e. 1.5× (or 34% savings) not only at the IP layer but also at the other layers. For a summary, see Table 5.12.

Sleep modes on a short time scale While the approaches discussed in the previous section exploit traffic variations on a daily time scale, there is additional energy saving potential on a packet-level time scale which can be partially independent from the previous approach. In packet switched networks, packet arrival rates can be highly non-uniform. In between packet arrivals, it is possible to off turn (sub)components. However, as the transition from operating to sleep state consumes both energy and time (which

---

[19] For example, in [87] it might not be instantly clear that the cost savings (which can act as a proxy for energy savings) listed in Table VII are for different overprovisioning values (maximum lightpath utilization $\delta = 1.0$ for LFA and GA, and maximum utilization of last lightpath on a logical link $\psi = 0.9$ for EWA) than the reference scenario (overprovisioning factor $\gamma = 0.5$).

might result in an increase in delay), there is a trade-off to the extent where this is possible. While there are several works on applying this approach in Ethernet LAN switches (e.g., Gupta [99], quoting savings in the range of 30% to 80%), we are aware of only the work by Nedevschi et al. [90] that considers it in the context of core networks, as shown in Table 5.11. They base their evaluations on packet-level simulations with real-world network topologies and traffic workloads. Their findings for the baseline of equipment where idle-mode power is almost comparable to active-mode power (as is the case for current telecom equipment) are savings in the range of 30% to 45% (IP layer savings, see Table 5.11). The above described approach is very similar to Energy-Efficient Ethernet (EEE), standardized as IEEE 802.3az, but which targets 100 Mbps to 10Gbps Ethernet copper-based physical layer devices. However, Reviriego et al. [100] found that already for link speeds of 1 Gbps and higher the proposed transition times (i.e., sleep and wakeup times, both in the order of microseconds[20]) are too high to achieve reasonable power proportionality. In addition, traffic patterns play a key role in the achievable savings. For optical transceivers, the transition times are even much larger (in the order of 1–2 milliseconds) than for copper-based devices due to a more complex circuitry to stabilize the channel, although recent work suggests that they can be dramatically reduced [91]. For high speeds optical links (40 Gbps and above), a promising alternative approach to circumvent the above issues with the wake/sleep transitions might be the use of a multilane architecture. That is, realizing parallel lower-rate lanes instead of one single high-rate lane; e.g. 4 × 10 Gbps to realize a single 40 Gbps link, not unlike the 'link bundling' concept discussed before. Depending on the traffic load, a number of lanes can be powered down and thus save energy. This has been studied in the context of energy efficiency by Revigiero et al. [92]; their results show indeed an improved power-load profile. So, while there are certainly some issues to resolve with respect to sleep modes at a packet-time scale, there do seem to be feasible solutions on the mid-term to realize important savings here. As the interdependency with sleep mode at daily time scale is not completely clear-cut, we assume for the **Moderate Effort reduction factor** 1.25× (or 20% savings) at the IP layer only, being a slightly lower value than the lower bound reported by Nedevschi et al. [90]. Similarly, for the **Best Effort reduction factor** we assume 1.4× (or 30% savings) for the IP layer and transponders (it seems unlikely that OXCs and OLAs can also benefit from this in the mid-term).

---

[20] For 10GBASE-T, the EEE wake and sleep transition times defined by IEEE 802.3az standard are 4.48 μs and 2.88 μs respectively, while the transmission time for a 1500 byte frame is 1.2 μs [100].

Table 5.12: *Reduction factor for the power rating factor*

	More eff. components[b]	Impr. chassis utilization	Sleep modes - daily	Sleep modes - short	(Sub)**total**	
Moderate Effort red.						
IP	*(1.15)*	1	1.15	1.25	1.44×	
Opt. Swit	-	1	1	1	1.00×	
Transponder	*(1.15)*	1	1	1	1.00×	
OLA	-	1	1	1	1.00×	
Weighted total[a]	*(1.13)*				**1.18**×	(=-15%)
Best Effort reduction						
IP	*(1.15)*	1.50	1.50	1.40	3.15×	
Opt. Swit	-	1.10	1.50	1	1.65×	
Transponder	*(1.15)*	1.10	1.50	1.40	2.31×	
OLA	-	1.50	1.50	1	2.25×	
Weighted total[a]	*(1.13)*				**2.62**×	(=-62%)

(a) See Section 5.3.3 for calculation details.

(b) The effect of more efficient components resulting from yearly efficiency improvements is *not* included in the weighted total in this table. See Section 5.5.1.2 for impact of yearly efficiency improvements.

Applying the proper weights to the values in Table 5.12, the overall result for the power rating P/C factor is a **Moderate Effort reduction factor** of 1.18× (or 15% savings), and a **Best Effort reduction factor** of 2.62× (or 62% savings), as shown in the same table.

5.4.6 Hop count H

The layer hop count H represents the average number of hops between processing elements in the respective layer, for example IP nodes in the IP layer. For a given network topology the hop count will depend on several aspects, such as the routing algorithm and link weights. However, a good ballpark number for H in a backbone network is 3–4 hops [1, 16, 21].

We consider the technique of optical bypass to reduce the hop count.

Table 5.13: Hop count related energy savings reported in publications

Source	Savings	Remarks	Justification
Optical Bypass			
Shen, 2009 [101]	15–45%	IP-over-WDM. MILP optimization and heuristics on three backbone networks (a six-node eight-link network, NSFNET and USNET). Upper range of savings for higher demands, and networks with more nodes.	Fig. 5 in [101], for the *'multi-hop bypass'* heuristic, which outperforms the more simple *'direct bypass'* heuristic.
Van Heddeghem, 2013 [96]	0–75%	IP-over-WDM. Shortest path routing across a variety of scenarios. Upper range of savings for sparse networks and high demand/linerate ratios.	Fig. 3–6 in [96].

Acronyms: MILP: Mixed-Integer Linear Programming.

An overview of the main works cited in the subsequent paragraph is given in Table 5.13.

Optical bypass A well-known technique to reduce the hop count H in the IP layer is to optically bypass IP routers, also known as IP offloading. The idea is that traffic not intended for the IP node remains in the optical layer and thus bypasses the IP router. The lightpath is switched, using OXCs, from an incoming fiber link directly on the appropriate outgoing fiber link. This allows us to reduce the capacity of the router and the associated power consumption. Optical bypass is possible at single-wavelength granularity, or on waveband granularity (requiring fewer ports in the OXC since multiple wavelengths are switched at the same time). The first thorough study to investigate optical bypass in the context of energy efficiency was by Shen et al. [101]. Next to an optimal upper bound on the savings, they evaluated two heuristics. While their 'direct bypass' strategy achieves savings around 5–45% (low–high demand), the 'multi-hop bypass' strategy (an intermediate, hybrid solution) was able to increase these savings with around 10 percentage points at lower demands.
From an energy (and also cost) perspective, optical bypass requires adequately filled optical channels, an issue not tackled in [101]. We have shown in [96] that if the lightpath filling ratio[21] is roughly below 50%, optical bypass consumes more energy than IP switching and grooming. Furthermore, in the same work [96] we've identified that another important parameter

---

[21] Note that lightpath filling relates to the concept of link bundles discussed in Section 5.4.5. For example, a link bundle consisting of 3 links would imply a lightpath filling ratio of 300%. However, the link bundle term cannot capture lightpath filling ratios $< 100\%$.

is the mesh degree of the network; full-mesh topologies benefit less from optical bypass.

As for small, full mesh core networks that serve traffic through high capacity interfaces (i.e., with demands relatively being small) the savings potential from employing optical bypass is limited (or even negative), we assume for the **Moderate Effort reduction factor** no savings, i.e. a reduction factor of 1.0×. On the other hand, if the average hop count in core networks is indeed around 3–4 hops, then there is considerable saving potential at the IP layer and transponders. Therefore, for the Best Effort scenario we have applied a factor 3 to the IP layer and transponders, as shown in Table 5.14. This results in a **Best Effort reduction factor** of 2.03× (or 51% savings), which is consistent with the results found in the above cited work. We should note that some PAR approaches from Section 5.4.5 might actually increase the hop count slightly, and therefore reduce the combined saving potential, which is one reason why we used the lower bound of hop count reductions at 3 instead of 4. Both reduction factors are summarized in Table 5.14.

Table 5.14: *Reduction factor for hop count*

	Optical bypass	(Sub)**total**	
Moderate Effort reduction	1	**1.00**×	(=-0%)
Best Effort reduction			
IP	3.00	3.00×	
Opt. Swit	1	1.00×	
Transponder	3.00	3.00×	
OLA	1	1.00×	
Weighted total[a]		**2.03**×	(=-51%)

[a] See Section 5.3.3 for calculation details.

5.4.7 Overprovisioning factor η_{op}

We have already mentioned the practice of overprovisioning before. Network operators typically overprovision the capacity of their backbone to account for peak-to-mean traffic variations, unexpected traffic spikes (e.g. in response to a major news event), failures and future traffic growth. We have separated the overprovisioning related to failures (i.e., protection) in

Table 5.15: *Overprovisioning factors reported in publications*

Source	Overpr.	Remarks
Zhang-Shen, 2004 [97]	10×	Given as '*. . . most networks today are enormously overprovisioned, with typical utilizations around 10%.*'
Huang, 2005 [102]	≥5×	Given as '*to accommodate the relatively large traffic fluctuations caused by link failures, values of $\beta \approx 5$ or even higher are not uncommon in large IP backbones*'
Fisher, 2010 [84]	2–3×	Given as '*The capacity of backbone networks is overprovisioned [. . .]. The average link utilization in backbone networks of large Internet service providers is estimated to be around 30–40%.*'
Kilper, 2011 [1]	2×	Does *not* include overprovisioning for protection.
Kilper, 2012 [4]	2–4×	Given as '*In the core network, the traffic is more uniform and the overprovisioning ratio can be as low as a factor of 2–4 times*', however this factor includes overprovisioning for protection as the author states '*overprovisioning is used to traffic bursts, provide spare capacity for use in the event of a failure (protection), and allow for mean traffic growth over a period of time*'

the protection factor η_{pr}, and dealt with it in Section 5.4.2. The reason we did so, is because there are numerous works that investigate specifically the energy saving potential of overprovisioning for protection, allowing a more accurate assessment of the power reduction potential.

In Table 5.15 we list the backbone overprovisioning factors as reported by a number of sources. From the values given, it is clear that there is a rather wide range of estimates on the actual overprovisioning factor. However, it is safe to assume that most of these factors should be halved, as they include the overprovisioning for protection; the work by Kilper et al. [1] is a notable exception. After doing this, the lower range of overprovisioning is a factor of 2.

The power consumption associated with this overprovisioned capacity can be reduced with the sleep mode techniques we discussed earlier in Section 5.4.5. For the Moderate Effort scenario we assume that at least a factor 2 can be saved at the IP layer, which results in a **Moderate Effort reduction factor** of 1.34× (or 25% savings). For the Best Effort scenario, we assume that a factor of 2 can be applied to all components in the network, resulting in a **Best Effort reduction factor** of 2× (or 50% savings). This is summarized in Table 5.16.

5.5 Total power saving potential and discussion

In this section we combine the power saving potential of the different approaches into a single number for the Moderate Effort and Best Effort sce-

Table 5.16: *Reduction factor for overprovisioning*

	Sleep modes	(Sub)**total**	
Moderate Effort reduction			
IP	2.00	2.00×	
Opt. Swit	1	1.00×	
Transponder	1	1.00×	
OLA	1	1.00×	
Weighted total(a)		**1.34**×	(=-25%)
Best Effort reduction	2.00	**2.00**×	(=-50%)

(a) See Section 5.3.3 for calculation details.

nario. We perform a sensitivity analysis on the two parameters in our model from Section 5.3 that could not be abstracted away, i.e., the average hop count H and the average link length. We compare our findings with the savings potential estimated by Greentouch's Green Meter [10]. Finally, we briefly assess how legacy networks (in contrast to our IP-over-WDM network) would impact our findings.

5.5.1 Total power saving potential

For discussing the impact of the reduction factors, we will distinguish between the static (once-only) approaches, and the time-dependent yearly efficiency improvements.

5.5.1.1 Static approaches only

In Fig. 5.9a we show the reduction factors for each of the static approaches, both in the Moderate Effort scenario (transparent bars) and the Best Effort scenario (filled bars)[22]. The time-dependent annual efficiency improvements are not included, but will be accounted for in the next section. The approaches with a high potential stand out clearly, i.e. reducing the overhead factor, the employment of optical bypass, and sleep modes for the overprovisioned capacity. However, to assess the realistic impact of each of the approaches, it makes sense to look at those that perform relatively

---

[22] The use of 2 decimal places for the reduction factors in Fig. 5.9a is to easily distinguish between the different factors, rather than being indicative for the level of accuracy of the factors.

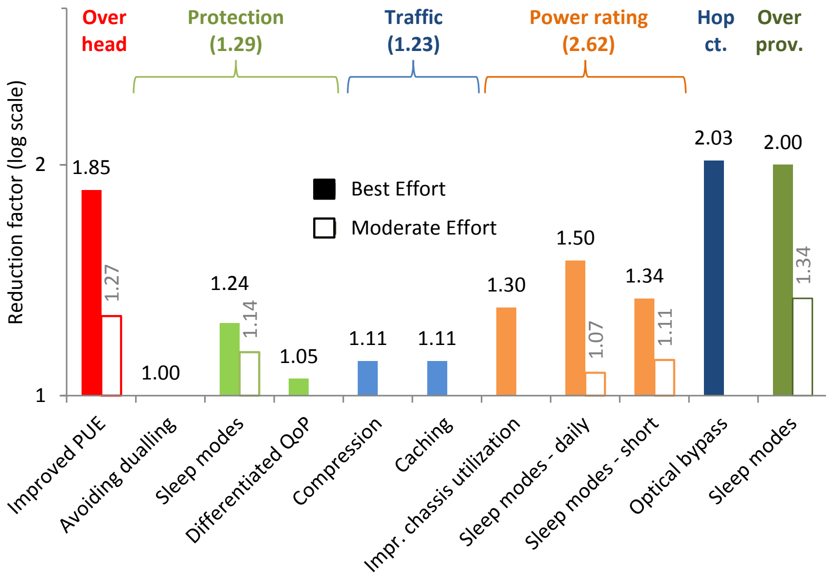

(a) Static reduction factors for each of the individual approaches (higher is better). The combined reduction potential of the power consumption in backbone networks is 2.3× (Moderate Effort) and 31× (Best Effort).

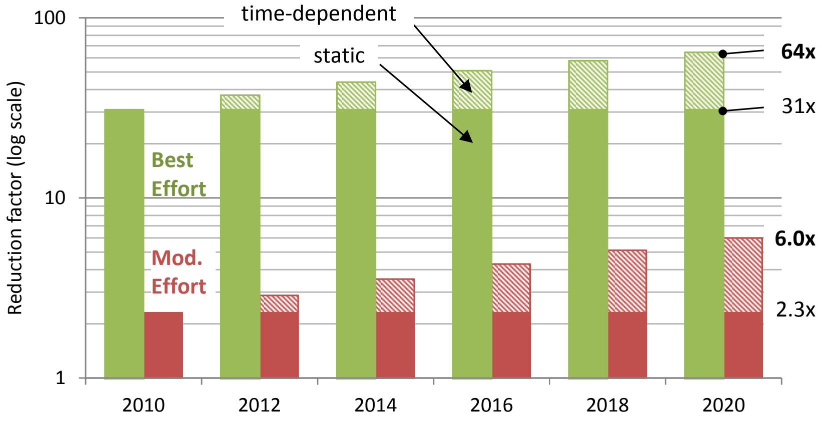

(b) Combined reduction factor potential of both the static approaches and the time-dependent 10% annual energy-per-bit reduction of telecom equipment (higher is better). By the year 2020, the combined reduction potential of the power consumption in backbone networks reaches 6.0× (Moderate Effort) and 64× (Best Effort).

Figure 5.9: *Reduction factors for the Moderate Effort and Best Effort scenarios for (a) the static-approaches only, and (b) the static and time-dependent approaches. Note that vertical axis uses a logarithmic scale so that the (sub)lengths are a measure for the individual reduction factors, and can easily be stacked.*

well in *both* the Moderate Effort and the Best Effort scenario. In this light, a reduction of power associated with the overhead factor and the overprovisioning, again, seems a promising and realistic direction. Another approach that performs well in both scenarios are sleep modes in general.

If we combine all of the static approaches depicted in Fig. 5.9a, not yet including the time-dependent annual efficiency improvement, the total reduction potential in the **Moderate Effort scenario** is **2.3**× (or 57% savings), and in the **Best Effort scenario 31**× (or 97% savings). Note that multiplying each of the reduction factors from Fig. 5.9a is a good approximation to get the combined reduction potential[23].

5.5.1.2 Static and time-dependent approaches

However, we have not yet accounted for the 11% yearly reductions in energy-per-bit (or 13% improvements in energy efficiency, which makes for a yearly reduction factor of 1.13), see Section 5.4.4. In Fig. 5.9b we plot for the time frame 2010 to 2020 the combined potential of both the static reduction factors (discussed above) and the impact of the expected yearly efficiency improvements. In both scenarios, the annual efficiency improvement is able to increase the static reduction potential with a little more[24] than a factor 2, resulting in a **Moderate Effort** reduction factor of **6.0**× (or 83% savings) and a **Best Effort** reduction factor of **64**× (or 98% savings).

For both calculations, the range between the Moderate Effort and Best Effort reduction factor is striking; they are more than an order of magnitude apart. The reason is twofold. First, the reported power saving potential for a similar approach can differ substantially between different works. Consider for example some of the values in Table 5.8 or in Table 5.11. This stems from differences in the assumed baseline (topologies, architectures, traffic matrices, power models), and the way in which power savings are evaluated and reported. Often maximum saving values are reported, at other times average savings, and sometimes it is plain unclear what the baseline is for the savings. This makes both the results potentially diverse, and very hard for the reader (and the authors of the current study) to compare energy saving results across different publications. As already noted

---

[23] Following our model in Section 5.3.3, the combined potential can strictly speaking *not* be calculated by multiplying each of the individual factors (as reduction factors differ for different equipment types, with equipment types having different weights). However, in our case the outcome of both the correct calculation as well as the simple multiplication approximation are nearly identical. For the Moderate Effort scenario we have a reduction factor of 2.30× (correct calculation) and 2.29× (approximation), for the Best Effort scenario we have 30.9× (correct calculation) and 31.4× (approximation).

[24] Again, we cannot simply multiply the static saving potential with $(1/87\%)^{10 years} = 4.03\times$, for the same reason as noted higher in the text.

by Bianzino et al. [103, 104], there is a clear need for a consistent methodology and benchmark to evaluate different power saving techniques. As of yet, this has not yet been sufficiently addressed. However, we do think that our survey at least partly addresses his suggestion for 'the creation [...] of a green repository of energy-related figures [...] to foster future research'. This survey should make it much more easier for authors to cross-check their results with community-wide results. The second reason for the wide range in both reduction factors is that part of our assessment is unavoidably subjective. We have tried to substantiate and justify why we chose the final reduction factors for each category as we did, and we did so to the best of our knowledge and the information we were able to find. Still, this is no 'hard science', and different authors would probably have made slightly different choices. However, we have made as much information available as we could, which should allow any reader to make his own assessment and estimation of the combined reduction factors.

5.5.2 Sensitivity to the average link length

As described in Section 5.3.2, for our estimations we assumed that backbone networks have an average link length of 800 km. This corresponds to each link having 10 OLAs. The end result is that OLAs contribute 9% to the total power consumption in the reference backbone network (see Table 5.2).

In Fig. 5.10 we show how the total static reduction factor is influenced by changing the link length to 300 km and 1300 km. As can be seen, higher average link lengths slightly decrease the power reduction factor. This is mainly driven by the contribution from applying optical bypass. When the link lengths increase, the relative weight of the OLAs to the total power consumption increases as well. As the OLAs don't contribute when applying optical bypass, the reductions of applying optical bypass are relatively decreased. For shorter link lengths, the inverse reasoning applies.

Overall, the influence of the average link length is relatively small, especially considering the level of accuracy associated with our reduction factors.

5.5.3 Sensitivity to the hop count H

The occurrence of the $(H + c_x)$ term in our model made our reduction factor calculation dependent on the average hop count H, see Section 5.3.3. In Fig. 5.11 we show the influence of changing the hop count $H = 3$ to either 2 or 5. It is important to note that we have changed the reduction potential of optical bypass to 2 and 5 respectively (i.e., its full potential), where it was 3 initially (see Table 5.14). As expected, an increase in the hop count increases

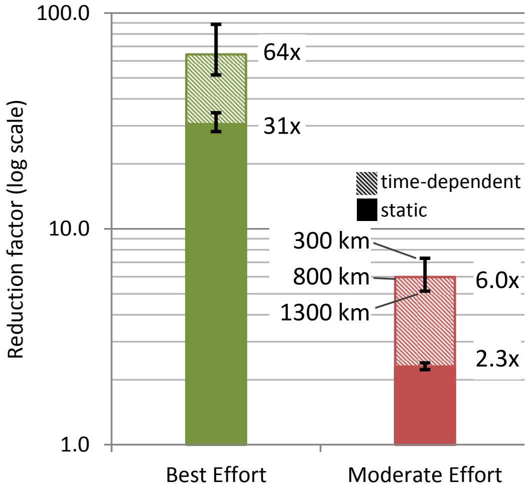

Figure 5.10: *Sensitivity of the combined reduction factor to the average link length (with reference network hop count $H = 3$). Higher average link lengths slightly decrease the power reduction potential. However, the sensitivity is relatively low.*

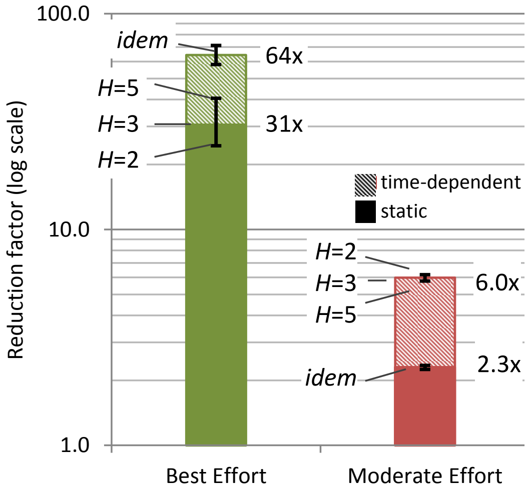

Figure 5.11: *Sensitivity of the combined reduction factor to the average hop count H (with reference network link length 800 km), over the range $H = 2 \rightarrow 5$. In practice, a higher hop count increases the reduction potential; this increase can be almost completely attributed to the fact that we also increased/decreased the power saving potential of employing optical bypass accordingly.*

Table 5.17: *GreenTouch's Green Meter power reduction estimations [10]*

Approach	Green Meter	Our study(a)
Improved components(b)	**27.0**×	-
PUE improvement (2.0 → 1.5)	*1.3*×	1.9×
Eff. impr. (13% p.a.)	*3.6*×	4.0×
Impr. optical interconnects	*1.5*×	-
Improved system design and integration	*5.0*×	-
Mixed line rates	**1.2**×	} 1.5×
Physical topology optim. to diurnal cycles	**1.1**×	
Optical bypass, sleep and low-energy state	**1.8**×	2.0×
Overall reduction factor	**64.0**×	-

[a] We list the Best Effort reduction factors. See text for more details.

[b] The net reduction potential (27×) is smaller than the 35.9× obtained by simply multiplying all the individual contributing factors. This is due to the same reason as in our study, i.e., the four individual factors apply only to the IP routers and transponders, with the PUE improvement also applying to the OXCs.

the total reduction factor, and vice versa. This can almost completely be attributed to the effect of employing optical bypass; for $H = 2$ the reduction potential of optical bypass drops to $1.56\times$, and for $H = 5$ goes up to $2.78\times$, compared to $2.03\times$ for $H = 3$.

If the reduction potential for optical bypass is not increased accordingly, e.g. as is the case for the Moderate Effort scenario (see Table 5.14), the effect of the hop count on the end result is negligible, as can be seen in Fig. 5.11.

5.5.4 Comparison with Green Meter

To our knowledge, the 'Green Meter v1.0'[25] white paper [10] published in July 2013 by the GreenTouch consortium is the only public report that adopts a similar methodology as we do in our study. Therefore it is especially relevant to compare our findings with the Green Meter findings.

There are several important things that the report brings to the table. First, the report does not exclusively focus on the core (backbone) network, but also includes the wireline access and mobile access network. In that sense it is a more holistic and complete approach (see Table 5.1 for an overview of the reported power reduction factors). Second, most of the different power reduction approaches are evaluated on a consistent benchmark framework, whereas most of the papers in our survey use different power values, reference networks and optimization parameters (the work by Idzikowski et al. [81] is a notable exception). On the other hand, our focus on the backbone network extends to research outside of GreenTouch.

---

[25] The Green Meter report is a work in progress, and subsequent updates and refinements to the model and estimations are planned.

This provides the reader with an insight in the consensus and uncertainty among the wider research community. Our report also provides both a Moderate Effort and a Best Effort scenario, which again contributes to an understanding of the range of achievable power savings. Finally, an independent approach allows to compare both works and perform a (limited) cross-validation, as we do here.

The reduction factors reported in the Green Meter study are listed in Table 5.17. The different approaches cannot exactly be mapped to our categorization, but nonetheless we can compare several approaches. The *PUE improvement factor* is listed as 1.3×. This is nearly identical to what we assume for the Moderate Effort scenario, whereas our Best Effort scenario goes as high as 1.85×. Our motivation for the latter is mainly grounded in the fact that the baseline PUE can be worse than 2.0. The 3.6× reduction factor resulting from the 13% of *yearly efficiency improvement* of IP routers and transponders is similar to our assumption of 15% efficiency improvement per year for the same equipment[26]. We do not specifically account for the *improved optical interconnects* and the *improved system design and integration*; if we would, they would fit under the power rating factor $\frac{P}{C}$ (most likely improved chassis utilization and sleep modes on a short time scale) and perhaps (partly) under the overprovisioning factor. Unfortunately, the Green Meter report does not provide details or references that underpin the very large reduction potential of improved system design and integration (5.0×), which makes it hard to properly assess its potential; part of this information is likely based on proprietary vendor research and not yet public. Likely it is an estimate of what clean-slate hardware design could bring to the table, where our focus is perhaps less aggressive and more about incremental improvements to hardware. The *mixed line rate* approach selects the power optimal combination of 40G, 100G and 400G interfaces for the given traffic distribution. We have not explicitly identified this as an approach, but we can consider it as similar to what we described as PAR (as PAR grooms traffic demands to more suitable-sized interfaces). However, we considered PAR in the context of exploiting daily traffic variations, which is similar to the Green Meter's *physical topology optimization to diurnal cycles*. Combining the savings of both approaches (1.2× and 1.1×) this gives a 1.3× reduction, which is not too far off of the 1.5× reduction we estimated for sleep modes that exploit daily traffic variations. Finally, Green Meter's *optical bypass* approach seems to include other techniques such as sleep modes and low-power states as well. Nonetheless, this maps rather well to our estimation of 2.03× for employing optical bypass.

---

[26] Don't confuse this with our overall backbone efficiency improvement of 13% per year, see Table 5.12.

Approaches that are not (yet) taken into account in the Green Meter, but that we do take into account are (a) protection related measures, (b) compression, and (c) caching.

Coincidentally, the Green Meter estimate of the combined potential (64×) is identical to our Best Effort reduction factor (64×). This is even more surprising, as we consider some approaches that Green Meter does not, and vice versa. The approach of the Green Meter seems to be more focussed on aggressive hardware optimisation and clean-slate design, whereas our focus is slightly more towards the application layers (cf. caching). On the other hand, our Moderate Effort reduction factor (4.9×) is well below the above figure.

5.5.5 The role of legacy networks

The baseline in our estimations is an IP-over-WDM network. This means we do not take into account the legacy network equipment and intermediate transport technologies such as Asynchronous Transfer Mode (ATM) and SDH that are probably[27] still widely present in current-day networks. As we already pointed out in Section 5.1, this is an important limitation and point to be aware of. It implies that our power reduction estimations cannot be interpreted as applicable to current-day networks. Rather, they should be seen as an estimation for the potential to improve current state-of-the art equipment.

Nonetheless, we can try to briefly assess how our estimations can be applied to legacy equipment and networks. From the approaches we identified in Fig. 5.9a, only few of them seem to be applicable to legacy equipment. Reducing the overhead power is certainly one of them, and fortunately provides good reduction potential (1.85×). Compression and caching probably can be applied to legacy networks as well, but the reduction potential (1.23×) is feeble and unsure. The other approaches would very likely require hardware upgrades to support features such as sleep modes, which is unrealistic for legacy equipment. And while putting equipment to sleep on a daily time scale might in theory be feasible, it is very unlikely that operators would resort to such measures with equipment that was not designed for it.

From our very brief assessment, it appears that if the power consumption of legacy network equipment is to be reduced drastically, the best option might be to replace it altogether with new, more energy-efficient equipment.

---

[27] We were not able to retrieve operator data on the actual distribution of deployed equipment; while part of this might be attributed to confidentiality issues, it seems not unlikely that up-to-date and network-wide inventories of such equipment are not readily available either.

5.6 Conclusion and recommendations

A novel methodology to categorize and survey power saving approaches In this paper we have surveyed a number of approaches to reduce the energy consumption of backbone IP-over-WDM telecommunication networks. The idea for this study was conceived when we noticed that in the available surveys the power saving potential (i.e. the reported saving percentages) of the works they survey was typically not discussed. This is striking, as estimating the energy saving potential of different techniques can be considered as one of the core drivers behind research in 'green networks'. The general approach in most existing surveys is a categorization of the various efficiency measures based on for example scale (e.g., optimizing at circuit-level vs. network level) or time (e.g., online sleep mode decisions vs. a preplanned energy-efficient network design). In contrast, our survey approach has been based on mapping the techniques onto a concise analytical power consumption model. This has the advantage that the power reduction potential associated with each of the different techniques can be rather easily multiplied in order to be combined to a total reduction factor.

Specifically, our power saving approaches have been disaggregated in those that reduce the power rating of equipment (a measure for their efficiency), the traffic in the network (i.e., the amount of transported bits), the number of network hops, the power associated with equipment installed for traffic protection purposes, equipment installed for dealing with unexpected traffic and future traffic growth (i.e., overprovisioning), and the power associated with external overhead such as cooling and power provisioning losses.

Power reduction potential in the backbone We assessed the power reduction potential for a Moderate Effort scenario and a Best Effort scenario, based on the uncertainty in the reported saving potential and (to a minor extent) operational and technical issues. We estimate the combined Moderate Effort reduction potential of once-only approaches at 2.3×, and their Best Effort reduction potential at 31×. If we factor in the projected yearly efficiency improvements driven by Moore's law (i.e. CMOS efficiency improvements), the 10-year reduction potential of both values are roughly (a little bit more than) doubled to 6.0× and 64× respectively. The large difference between the Moderate Effort and Best Effort estimate is caused mainly by the disparity and the lack of clarity of the reported savings, and our (to an extent unavoidable, subjective) assessment of the feasibility of implementation. The large difference also reflects the uncertainty we encountered when evaluating many of the surveyed works, and confirms earlier conclusions by Bianzino et al. that a standardized methodology and benchmark is deeply needed. Nonetheless we think that the general trends and findings are sound. We do provide ample information in our survey for the reader to make his own assessment.

On an individual level, a number of approaches stand out with their reduction potential. We found that the approaches that focus on bypassing power-hungry IP hops, reducing the power associated with the external overhead (i.e. improving the PUE), and putting overcapacity to sleep can reduce the power consumption most; each with roughly a factor 2 in the Best Effort scenario. Especially the latter two techniques also provide a relatively high Moderate Effort reduction potential, which is an indication for our confidence in their potential. More general, the technique of applying sleep modes shows good potential across a variety of applications.

Outrunning the traffic growth? Considering that the traffic is projected to increase close to 10-fold in the next 10 years, the power reduction potential of our Moderate Effort scenario (6.0×) will not be sufficient to halt the power consumption increase in backbone telecommunication networks. However, the estimated 31–64× reduction for our Best Effort scenario shows that enough potential is likely available to halt, and even reduce, the absolute power consumption. This will no doubt require significant efforts from an economic, technical and operational point of view. These issues have not been considered in our survey, and require extensive investigations on their own.

Recommendations for future research From what we learned while doing this survey, we would like to make the following recommendations to our readers. For the research community: (a) Any paper on energy-efficient network techniques should clearly distinguish the various effects that result in a power consumption reduction; e.g., side effects such as an increase in the link load should be clearly identified. (b) The baseline scenario should be clearly identified; often we found that it was ambiguous whether the baseline was e.g. a common shortest-path, minimized delay scenario, or another energy-optimized scenario. (c) It should be clearly stated whether the reported saving potential applies to average or maximum savings, and to which equipment it applies (e.g., only line cards, IP equipment, or the complete network). (d) There is a need for studies that compare various power saving approaches across a consistent baseline (such as topology, power model, and traffic model). This allows for a cross-validation of reported results.

For network operators, vendors and policy makers alike, we would like to state that there is really a need for a more open spirit with respect to the publication of power consumption values of network devices and complete networks. While research into energy efficiency of ICT is already several years ongoing, still very little credible power consumption values from vendors are available and a large range of equipment power consumption values are used in various works. The availability of such values directly from vendors would probably improve the consistency of the power consumption values used as an input in future research.

Finally, we hope that this survey is useful for researchers in the short-term to mid-term to check and validate their own (intermediate) power saving results across a wide set of existing works, and thereby perform a first validation of their findings.

Acknowledgments

The authors would like to thank the following people for their feedback and input: Gianluca Griffa, Peter Lambert, Sofie Lambert, Christoph Lange, Esther Le Rouzic, and Andrew Lord. Of course, any errors or misinterpretations are the responsibility of the authors. The work described in this paper was carried out with the support of the iMinds-project GreenICT, and the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257740 (Network of Excellence TREND).

References

[1] D. C. Kilper, G. Atkinson, S. K. Korotky, S. Goyal, P. Vetter, D. Suvakovic, and O. Blume. *Power Trends in Communication Networks*. IEEE Journal of Selected Topics in Quantum Electronics, 17(2):275–284, March 2011. doi:10.1109/JSTQE.2010.2074187.

[2] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch. *Energy Consumption of Telecommunication Networks and Related Improvement Options*. Selected Topics in Quantum Electronics, IEEE Journal of, 17(2):285–295, April 2011. doi:10.1109/JSTQE.2010.2053522.

[3] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker. *Energy Consumption in Optical IP Networks*. Journal of Lightwave Technology, 27(13):2391–2403, July 2009. doi:10.1109/JLT.2008.2010142. Available from: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4815495.

[4] D. Kilper, K. Guan, K. Hinton, and R. Ayre. *Energy Challenges in Current and Future Optical Transmission Networks*. Proceedings of the IEEE, 100(5):1168–1187, May 2012. doi:10.1109/JPROC.2012.2186102.

[5] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee. *Energy Efficiency in Telecom Optical Networks*. Communications Surveys & Tutorials, IEEE, 12(4):441–458, November 2010. doi:10.1109/SURV.2011.062410.00034.

[6] R. Bolla, R. Bruschi, F. Davoli, and F. 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, May 2011. doi:10.1109/SURV.2011.071410.00073.

[7] A. P. Bianzino, C. Chaudet, D. Rossi, and J.-L. Rougier. *A Survey of Green Networking Research*. Communications Surveys & Tutorials, IEEE, 14(1):3–20, February 2012. doi:10.1109/SURV.2011.113010.00106.

[8] S. Roy. *Energy Logic for Telecommunications*. Emerson, White Paper, 2008.

[9] M. Parker and S. Walker. *Roadmapping ICT: an Absolute Energy Efficiency Metric*. Optical Communications and Networking, IEEE/OSA Journal of, 3(8):A49–A58, July 2011. doi:10.1364/JOCN.3.000A49.

[10] GreenTouch. *GreenTouch Green Meter Research Study: Reducing the Net Energy Consumption in Communications Networks by up to 90% by 2020*. http://www.greentouch.org/uploads/documents/GreenTouch_Green_Meter_Research_Study_26_June_2013.pdf, June 2013.

[11] E. Le Rouzic, E. Bonetto, L. Chiaraviglio, F. Giroire, F. Idzikowski, F. Jimenez, C. Lange, J. Montalvo, F. Musumeci, I. Tahiri, A. Valenti, W. Van Heddeghem, Y. Ye, A. Bianco, and A. Pattavina. *TREND Towards More Energy-efficient Optical Networks*. In ONDM, pages 211–216, Brest, France, April 2013.

[12] E. Le Rouzic, R.-M. Indre, L. Chiaraviglio, F. Musumeci, A. Pattavina, J. Lopez Vizcaino, Y. Ye, W. Van Heddeghem, A. Bianco, E. Bonetto, M. Meo, F. Jimenez, F. Idzikowski, and R. Cuevas. *TREND Big Picture on Energy-efficient Backbone Networks*. In Digital Communications-Green ICT (TIWDC), 2013 24th Tyrrhenian International Workshop on, pages 1–6. IEEE, September 2013. doi:10.1109/TIWDC.2013.6664209.

[13] W. Van Heddeghem, M. C. Parker, S. Lambert, W. Vereecken, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester. *Using an Analytical Power Model to Survey Power Saving Approaches in Backbone Networks*. In Networks and Optical Communications (NOC), 2012 17th European Conference on, pages 1–6, Vilanova i la Geltru, Spain, June 2012. IEEE. doi:10.1109/NOC.2012.6249942.

[14] J. Malmodin, A. Moberg, D. Lundén, G. Finnveden, and N. Lövehagen. *Greenhouse Gas Emissions and Operational Electricity Use in the ICT, and Entertainment & Media Sectors*. J. Ind. Ecol., 14(5):770–790, October 2010. doi:10.1111/j.1530-9290.2010.00278.x.

[15] 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, December 2012. doi:10.1364/OE.20.00B513.

[16] W. Van Heddeghem, F. Idzikowski, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester. *Power Consumption Modeling in Optical Multilayer Networks*. Photonic Network Communications, January 2012. doi:10.1007/s11107-011-0370-7. Available from: http://www.springerlink.com/index/10.1007/s11107-011-0370-7.

[17] J. C. C. Restrepo, C. G. Gruber, and C. M. Machuca. *Energy Profile Aware Routing*. In Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on, pages 1–5, Dresden, Germany, June 2009. IEEE. doi:10.1109/ICCW.2009.5208041.

[18] R. Bolla, R. Bruschi, C. Lombardo, and D. Suino. *Evaluating the Energy-awareness of Future Internet Devices*. In High Performance Switching and Routing (HPSR), 2011 IEEE 12th International Conference on, pages 36 –43, July 2011. doi:10.1109/HPSR.2011.5986001.

[19] E. R. Masanet, R. E. Brown, A. Shehabi, J. G. Koomey, and B. Nordman. *Estimating the Energy Use and Efficiency Potential of US Data Centers*. Proceedings of the IEEE, 99(8):1440–1453, August 2011. doi:10.1109/JPROC.2011.2155610.

[20] W. Van Heddeghem, F. Idzikowski, E. Le Rouzic, J. Mazeas, H. Poignant, S. Salaun, B. Lannoo, and D. Colle. *Evaluation of Power Rating of Core Network Equipment in Practical Deployments*. In Online Conference on Green Communications (GreenCom), 2012 IEEE, pages 126–132, September 2012. doi:10.1109/GreenCom.2012.6519628.

[21] R. Hülsermann, A. Betker, M. Jäger, S. Bodamer, M. Barry, J. Späth, C. Gauger, and M. Köhn. *A Set of Typical Transport Network Scenarios for Network Modelling*. ITG FACHBERICHT, (182):65–72, 2004.

[22] C. Belady, A. Rawson, J. Pfleuger, and T. Cader. *Green Grid Data Center Power Efficiency Metrics: PUE and DCIE*. White Paper #6, The Green Grid, 2008.

[23] V. Avelar, D. Azevedo, and A. French. *PUE: A Comprehensive Examination of the Metric*. White Paper #49, The Green Grid, 2012.

[24] EPA. *Report to Congress on Server and Data Center Energy Efficiency*. U.S. Environmental Protection Agency, http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf, 2007.

[25] Google. *Data Center Efficiency – Power Usage Effectiveness*. http://www.google.com/about/datacenters/inside/efficiency/power-usage.html, 2012.

[26] M. Stansberry and J. Kudritzki. *Uptime Institute 2012 Data Center Industry Survey*. Uptime Institute, 2012. Available at: http://uptimeinstitute.com/2012-survey-results.

[27] 42u. *What is PUE? What is DCiE?* http://www.42u.com/measurement/pue-dcie.htm#What-is-PUE?

[28] N. Rasmussen. *Implementing Energy Efficient Data Centers (White Paper 114, Revision 1)*. APC, 2011.

[29] Intel. *The State of Data Center Cooling - A Review of Current Air and Liquid Cooling Solutions*. ftp://download.intel.com/technology/EEP/data-center-efficiency/state-of-date-center-cooling.pdf, 2008.

[30] Google. *Data Center Best Practices*. http://www.google.com/about/datacenters/best-practices.html, 2012.

[31] 42u. *Cooling Solutions for Server Racks, Computer Rooms, Server Rooms and Data Centers*. http://www.42u.com/42u-rack-cooling.htm.

[32] E. Pinheiro, W. Weber, and L. Barroso. *Failure Trends in a Large Disk Drive Population*. In Proceedings of the 5th USENIX conference on File and Storage Technologies, pages 2–2, 2007.

[33] S. Greenberg, E. Mills, B. Tschudi, P. Rumsey, and B. Myatt. *Best Practices for Data Centers: Lessons Learned from Benchmarking 22 Data Centers*. Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings in Asilomar, CA. ACEEE, 3:76–87, August 2006.

[34] X. Fan, W. Weber, and L. Barroso. *Power Provisioning For A Warehouse-sized Computer*. ACM SIGARCH Computer Architecture News, 35(2):13–23, May 2007. doi:10.1145/1273440.1250665.

[35] J. Koomey. *The NYT article on Power, Pollution, and the Internet: My initial Comments*, September 2012. Blog post available at: http://www.koomey.com/post/32281701993.

[36] W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester. *Trends in Worldwide ICT Electricity Consumption from 2007 to 2012*. Computer Communications, 50:64–76, September 2014. doi:10.1016/j.comcom.2014.02.008.

[37] *Responsabilité sociale d'entreprise, rapport détaillé*. France Telecom - Orange, 2011. Available at: http://www.orange.com/fr/finance/nbsp4/assemblee-generale-2013/rapports-annuels.

[38] P. Blasco. *Telefónica inaugura en Alcalá de Henares uno de los mayores centros de datos [Telefonica opens in Alcala de Henares one of the largest data centers]*. El Mundo.es, April 30, 2013. Available at: http://www.elmundo.es/elmundo/2013/04/30/navegante/1367304168.html.

[39] *The 2010/2011 Corporate Responsibility Report*. Deutsche Telekom. Retrieved on 2013-11-28 at: http://www.cr-report.telekom.com/site11/en/klima-und-umwelt/umweltfreundliche-netzinfrastruktur/klimatechnik-und-energieeinsatz/index.php?page=1#ff,43918,43917,43916,43915,43914,43913,43920,43919.

[40] *Corporate Responsibility Report 2012*. Deutsche Telekom. Retrieved on 2013-11-28 at: http://www.cr-report.telekom.com/site13/sites/default/files/pdf/2012_corporate_responsibility_report-individuell.pdf.

[41] S. Le Masson and D. Nortershauser. *Complex Wall for Indirect Freecooling in Datacentres*. In Telecommunications Energy Conference (INTELEC), 2012 IEEE 34th International, pages 1–6, Scottsdale, USA, October 2012. IEEE. doi:10.1109/INTLEC.2012.6374506.

[42] C. Cavdar, F. Buzluca, and L. Wosinska. *Energy-efficient Design of Survivable WDM Networks with Shared Backup*. In Global Telecommunications Conference (GLOBECOM 2010), 2010 IEEE, pages 1–5, Miami, USA, December 2010. IEEE. doi:10.1109/GLOCOM.2010.5683203.

[43] A. Jirattigalachote, C. Cavdar, P. Monti, L. Wosinska, and A. Tzanakaki. *Dynamic Provisioning Strategies for Energy Efficient WDM Networks with Dedicated Path Protection*. Optical Switching and Networking, 8(3):201–213, July 2011. doi:10.1016/j.osn.2011.03.008. Available from: http://linkinghub.elsevier.com/retrieve/pii/S1573427711000208.

[44] A. Coiro, M. Listanti, A. Valenti, and F. Matera. *Reducing Power Consumption in Wavelength Routed Networks by Selective Switch Off of Optical Links*. Selected Topics in Quantum Electronics, IEEE Journal of, 17(2):428–436, April 2011. doi:10.1109/JSTQE.2010.2052791.

[45] N.-H. Bao, L.-M. Li, H.-F. Yu, Z.-Z. Zhang, and H.-B. Luo. *Power-aware Provisioning Strategy with Shared Path Protection in Optical WDM Networks*. Optical Fiber Technology, 18(2):81–87, March 2012. doi:10.1016/j.yofte.2011.12.003.

[46] F. Musumeci, M. Tornatore, J. Lopez Vizcaino, Y. Ye, and A. Pattavina. *Energy-efficiency of protected IP-over-WDM networks with sleep-mode devices*. Journal of High Speed Networks, 19(1):19–32, March 2013. doi:10.3233/JHS-130463.

[47] J. Lopez Vizcaino, Y. Ye, V. Lopez, F. Jimenez, R. Duque, F. Musumeci, A. Pattavina, and P. Krummrich. *Differentiated Quality of Protection*

to Improve Energy Efficiency of Survivable Optical Transport Networks. In Optical Fiber Communication Conference, Anaheim, USA, March 2013. Optical Society of America. doi:10.1364/OFC.2013.OM3A.1.

[48] O. Tamm, C. Hermsmeyer, and A. M. Rush. *Eco-sustainable System and Network Architectures for Future Transport Networks*. Bell Labs Technical Journal, 14(4):311–327, Winter 2010. doi:10.1002/bltj.20418.

[49] X. Dong, A. Lawey, T. E. El-Gorashi, and J. M. Elmirghani. *Energy-efficient Core Networks*. In Optical Network Design and Modeling (ONDM), 2012 16th International Conference on, pages 1–9, Colchester, UK, April 2012. IEEE. doi:10.1109/ONDM.2012.6210196.

[50] R. Kothiyal, V. Tarasov, P. Sehgal, and E. Zadok. *Energy and Performance Evaluation of Lossless File Data Compression on Server Systems*. In Proceedings of SYSTOR 2009: The Israeli Experimental Systems Conference, page 4, Haifa, Israel, May 2009. ACM. doi:10.1145/1534530.1534536.

[51] U. Lee, I. Rimac, and V. Hilt. *Greening the Internet with Content-centric Networking*. In Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking (e-Energy '10), pages 179–182, Passau, Germany, April 2010. ACM. doi:10.1145/1791314.1791342.

[52] U. Lee, I. Rimac, D. Kilper, and V. Hilt. *Toward Energy-efficient Content Dissemination*. Network, IEEE, 25(2):14–19, April 2011. doi:10.1109/MNET.2011.5730523.

[53] N. Osman, T. El-Gorashi, and J. Elmirghani. *Reduction of Energy Consumption of Video-on-demand Services Using Cache Size Optimization*. In Wireless and Optical Communications Networks (WOCN), 2011 Eighth International Conference on, pages 1–5, Paris, France, May 2011. IEEE. doi:10.1109/WOCN.2011.5872923.

[54] N. I. Osman, T. El-Gorashi, and J. M. Elmirghani. *Caching in Green IP over WDM networks*. Journal of High Speed Networks, 19(1):33–53, March 2013. doi:10.3233/JHS-120458.

[55] L. Chiaraviglio and I. Matta. *An Energy-aware Distributed Approach for Content and Network Management*. In Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on, pages 337–342, Shanghai, China, April 2011. IEEE. doi:10.1109/INFCOMW.2011.5928834.

[56] L. Chiaraviglio and I. Matta. *Energy-Aware Network Management and Content Distribution*. In J. Wu, S. Rangan, and H. Zhang, editors, Green Communications: Theoretical Fundamentals, Algorithms, and Applications. CRC Press, 2012.

[57] R. Modrzejewski, L. Chiaraviglio, I. Tahiri, F. Giroire, E. Le Rouzic, E. Bonetto, F. Musumeci, R. Gonzalez, and C. Guerrero. *Energy Efficient Content Distribution in an ISP Network*. In Global Communications Conference (GLOBECOM), 2013 IEEE, Atlanta, USA, December 2013. doi:10.1109/GLOCOM.2013.6831508.

[58] U. Mandal, M. F. Habib, S. Zhang, C. Lange, A. Gladisch, and B. Mukherjee. *Adopting Hybrid CDN–P2P in IP-over-WDM Networks: an Energy-efficiency Perspective*. Journal of Optical Communications and Networking, 6(3):303–314, February 2014. doi:10.1364/JOCN.6.000303.

[59] Cisco. *Cisco Visual Networking Index: Forecast and Methodology, 2010–2015*, 2011.

[60] Cisco. *Global IP Traffic Forecast and Methodology, 2006–2011*, 2007.

[61] Cisco. *Cisco Visual Networking Index: Forecast and Methodology, 2008–2013*, 2009.

[62] Cisco. *Cisco Visual Networking Index: Forecast and Methodology, 2012–2017*, 2013.

[63] Y. Wu, P. A. Chou, and S.-Y. Kung. *Minimum-energy Multicast in Mobile Ad Hoc Networks Using Network Coding*. Communications, IEEE Transactions on, 53(11):1906–1918, November 2005. doi:10.1109/TCOMM.2005.857148.

[64] X. Liu, K. Fouli, R. Kang, and M. Maier. *Network-coding-based Energy Management for Next-generation Passive Optical Networks*. Journal of Lightwave Technology, 30(6):864–875, March 2012. doi:10.1109/JLT.2011.2182633.

[65] GreenTouch. *GreenTouch 2012–2013 Annual Report*. http://www.greentouch.org/uploads/documents/GreenTouch_2013_Annual_Report.pdf, June 2013.

[66] J. Koomey, S. Berard, M. Sanchez, and H. Wong. *Implications of Historical Trends in The Electrical Efficiency of Computing*. Annals of the History of Computing, IEEE, 33(3):46–54, March 2011. doi:10.1109/MAHC.2010.28.

[67] T. Ebrahimi. *What Future for Image/Video Compression*. http://www.icme2011.org/files/icme2011_panel1_ebrahimi.pdf, July 2011.

[68] G. Heydon. *IPTV (is changing fast)*. ITU-AIBD-ABU Regional Workshop – Strengthening Digital Broadcasting Experience, Hanoi (Vietnam), available at: http://www.itu.int/ITU-D/asp/CMS/Events/2011/DigitalBroadcast-May2011/S8_GeofHeydon.pdf, May 2011.

[69] G. Peng. *CDN: Content Distribution Network*. Arxiv preprint cs/0411069, 2004.

[70] A. Feldmann, A. Gladisch, M. Kind, C. Lange, G. Smaragdakis, and F. Westphal. *Energy Trade-offs Among Content Delivery Architectures*. In Telecommunications Internet and Media Techno Economics (CTTE), 2010 9th Conference on, pages 1–6, Ghent, Belgium, June 2010. IEEE. doi:10.1109/CTTE.2010.5557700.

[71] D. T. Neilson. *Photonics for Switching and Routing*. Selected topics in quantum electronics, IEEE Journal of, 12(4):669–678, August 2006. doi:10.1109/JSTQE.2006.876315.

[72] S. Han. *Moore's Law and Energy and Operations Savings in the Evolution of Optical Transport Platforms*. Communications Magazine, IEEE, 48(2):66–69, February 2010. doi:10.1109/MCOM.2010.5402665.

[73] L. Ceuppens. *Planning for Energy Efficiency: Networking in Numbers*. In OFC/NFOEC 2009, San Diego, USA, 2009.

[74] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright. *Power Awareness in Network Design and Routing*. In Proceedings of INFOCOM, pages 1130–1138, Phoenix, USA, April 2008. doi:10.1109/INFOCOM.2008.93.

[75] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan. *A Power Benchmarking Framework for Network Devices*. In Proceedings of NETWORKING, Aachen, Heidelberg, pages 795–808, Aachen, Germany, 2009. doi:10.1007/978-3-642-01399-7_62.

[76] R. Tucker. *Green Optical Communications – part II: Energy Limitations in Networks*. Selected Topics in Quantum Electronics, IEEE Journal of, 17(2):261–274, April 2011. doi:10.1109/JSTQE.2010.2051217.

[77] M. C. Parker, S. D. Walker, R. Martin, W. Van Heddeghem, and B. Lannoo. *Energy-Efficient Master-Slave Edge-Router Upgrade Paths in*

Active Remote Nodes of Next-Generation Optical Access. In Asia Communications and Photonics Conference, Guangzhou, China, November 2012. Optical Society of America.

[78] C. Lange and A. Gladisch. *Energy Efficiency Limits of Load Adaptive Networks*. In Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference, 2010 Conference on (OFC/NFOEC), pages 1–3, San Diego, USA, March 2010. IEEE. doi:10.1364/OFC.2010.OWY2.

[79] L. Liu and B. Ramamurthy. *Rightsizing Bundle Link Capacities for Energy Savings in the Core Network*. In Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE, pages 1–6, Houston, USA, December 2011. IEEE. doi:10.1109/GLOCOM.2011.6134490.

[80] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz. *Dynamic Routing at Different Layers in IP-over-WDM Networks - Maximizing Energy Savings*. Optical Switching and Networking, 8(3):181–200, July 2011. doi:10.1016/j.osn.2011.03.007.

[81] F. Idzikowski, E. Bonetto, L. Chiaraviglio, A. Cianfrani, A. Coiro, R. Duque, F. Jiménez, E. Le Rouzic, F. Musumeci, W. Van Heddeghem, J. López Vizcaíno, and Y. Ye. *TREND in Energy-Aware Adaptive Routing Solutions*. Communications Magazine, IEEE, November 2013. doi:10.1109/MCOM.2013.6658659.

[82] W. Vereecken, W. Van Heddeghem, M. Deruyck, B. Puype, B. Lannoo, W. Joseph, D. Colle, L. Martens, and P. Demeester. *Power Consumption in Telecommunication Networks: Overview and Reduction Strategies*. Communications Magazine, IEEE, 49(6):62–69, June 2011. doi:10.1109/MCOM.2011.5783986.

[83] Y. Zhang, M. Tornatore, P. Chowdhury, and B. Mukherjee. *Time-aware Energy Conservation in IP-over-WDM Networks*. In Photonics in Switching. Optical Society of America, July 2010. doi:10.1364/PS.2010.PTuB2.

[84] W. Fisher, M. Suchara, and J. Rexford. *Greening Backbone Networks: Reducing Energy Consumption by Shutting Off Cables in Bundled Links*. In Proceedings of the first ACM SIGCOMM workshop on Green networking, pages 29–34. ACM, 2010. doi:10.1145/1851290.1851297.

[85] F. Idzikowski, L. Chiaraviglio, R. Duque, F. Jiménez, and E. Le Rouzic. *Green Horizon: Looking at Backbone Networks in 2020 from the Perspective of Network Operators*. In Communications (ICC), 2013 IEEE Inter-

national Conference on, pages 4455–4460, Budapest, Hungary, June 2013. IEEE. doi:10.1109/ICC.2013.6655268.

[86] L. Chiaraviglio, M. Mellia, and F. Neri. *Minimizing ISP Network Energy Cost: Formulation and Solutions*. IEEE/ACM Transactions on Networking (TON), 20(2):463–476, April 2012. doi:10.1109/TNET.2011.2161487.

[87] E. Bonetto, L. Chiaraviglio, F. Idzikowski, and E. Le Rouzic. *Algorithms for the Multi-Period Power-Aware Logical Topology Design With Reconfiguration Costs*. Journal of Optical Communications and Networking, 5(5):394–410, April 2013. doi:10.1364/JOCN.5.000394.

[88] R. Bolla, R. Bruschi, A. Cianfrani, and M. Listanti. *Introducing Standby Capabilities into Next-generation Network Devices*. In Proceedings of the PRESTO, Philadelphia, USA, November 2010.

[89] A. Coiro, M. Listanti, A. Valenti, and F. Matera. *Energy-aware Traffic Engineering: a Routing-based Distributed Solution for Connection-oriented IP Networks*. Computer Networks, 57(9):2004–2020, June 2013. doi:10.1016/j.comnet.2013.03.017.

[90] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall. *Reducing Network Energy Consumption via Sleeping and Rate-adaptation*. In Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation, 2008.

[91] P. Reviriego, D. Larrabeiti, and J. A. Maestro. *Would Energy Efficient Ethernet be effective on 10Gbps Optical Links?* In Photonics in Switching. Optical Society of America, July 2010. doi:10.1364/PS.2010.PTuB5.

[92] P. Reviriego, B. Huiszoon, V. López, R. Coenen, J. Hernández, and J. Maestro. *Improving Energy Efficiency in IEEE 802.3 ba High-Rate Ethernet Optical Links*. Selected Topics in Quantum Electronics, IEEE Journal of, 17(2):419–427, April 2011. doi:10.1109/JSTQE.2010.2050136.

[93] F. Ilchmann et al. *Efficient and Optimized Network Architecture: Requirements and Reference Scenarios*. Deliverable D2.1, STRONGEST Project, September 2010.

[94] Labovitz, Craig. *What Europeans do at Night*. ARBOR Networks, blog post at http://www.arbornetworks.com/asert/2009/08/what-europeans-do-at-night/, August 2009.

[95] R. Doverspike, K. K. Ramakrishnan, and C. Chase. *Structural Overview of ISP Networks*. In C. Kalmanek, S. Misra, and R. Yang, editors, Guide to Reliable Internet Services and Applications, pages 19–93. Springer, 2010. doi:10.1007/978-1-84882-828-5_2.

[96] W. Van Heddeghem, F. Musumeci, F. Idzikowski, A. Pattavina, B. Lannoo, D. Colle, and M. Pickavet. *Power Consumption Evaluation of Circuit-Switched Versus Packet-Switched Optical Backbone Networks*. In Online Conference on Green Communications (GreenCom), 2013 IEEE, pages 56–63, October 2013. doi:10.1109/OnlineGreenCom.2013.6731029.

[97] R. Zhang-Shen and N. McKeown. *Designing a Predictable Internet Backbone Network*. In Proceedings of HotNets III, San Diego, USA, 2004.

[98] E. Tego, F. Idzikowski, L. Chiaraviglio, A. Coiro, and F. Matera. *Facing the Reality: Validation of Energy Saving Mechanisms on a Testbed*. Journal of Electrical and Computer Engineering, 2014, 2014. doi:10.1155/2014/806960.

[99] M. Gupta and S. Singh. *Using Low-Power Modes for Energy Conservation in Ethernet LANs*. In INFOCOM, volume 7, pages 2451–55, 2007.

[100] P. Reviriego, J. Hernández, D. Larrabeiti, and J. A. Maestro. *Performance Evaluation of Energy Efficient Ethernet*. Communications Letters, IEEE, 13(9):697–699, September 2009. doi:10.1109/LCOMM.2009.090880.

[101] G. Shen and R. Tucker. *Energy-minimized Design for IP over WDM networks*. Journal of Optical Communications and Networking, 1(1):176–186, June 2009. doi:10.1364/JOCN.1.000176.

[102] Y. Huang and R. Guerin. *Does Over-provisioning Become More or Less Efficient as Networks Grow Larger?* In Network Protocols, 2005. ICNP 2005. 13th IEEE International Conference on, pages 11–pp, Boston, USA, November 2005. IEEE. doi:10.1109/ICNP.2005.12.

[103] A. P. Bianzino, A. K. Raju, and D. Rossi. *Apples-to-apples: a Framework Analysis for Energy-efficiency in Networks*. SIGMETRICS Perform. Eval. Rev., 38(3):81–85, January 2011. doi:10.1145/1925019.1925036. Available from: http://doi.acm.org/10.1145/1925019.1925036.

[104] A. P. Bianzino, A. Raju, and D. Rossi. *Power Profiling the Internet Core: A Sensitivity Analysis*. Energy-aware Systems and Networking for

Sustainable Initiatives, page 287, 2012. doi:10.4018/978-1-4666-1842-8.ch013.

6

Conclusion

"Essentially, all models are wrong, but some are useful."

–George E.P. Box (1919–2013)

The starting point for this dissertation was the observation that the electricity consumption of Information and Communication Technology (ICT) was growing at 8% per year around 2008. This effectively means that its electricity consumption doubles every 9 years, which is clearly not sustainable.

The power consumption of ICT As a first contribution in this dissertation we wanted to assess whether the yearly growth in ICT electricity consumption has changed since the 2008 study, given the increased worldwide attention for energy-efficiency in all sectors, including ICT. We estimated the electricity consumption of communication networks, personal computers and data centers in 2007 and 2012. Our results indicate that the combined growth across these three categories is now 7% per year; this is lower than the 10% per year observed before 2007 for the same three categories (the 2008 study mentioned above also considered TVs and an 'Other ICT equipment' category, which explains the difference in growth rates mentioned here). So indeed, **ICT electricity consumption is now growing at a rate which is smaller than before 2007**. An important reason for this slowdown is the shift to more efficient technologies: from desktops to laptops, from bulky CRT monitors to LCD monitors, implementing server virtualization and more efficient cooling in data centers. **While this is in a sense good news, we should not be blind for the fact that it is still increasing faster than the global human electricity consumption**. All three categories—communication networks, personal computers, and data centers—consume roughly an equal amount of energy. The highest growth rates are observed in telecommunication networks, which is not surprising given the explosive growth of mobile communication in the last decade.

Reducing the energy consumption in backbone networks The observations above make a strong case for intensifying the efforts to reduce the ICT energy consumption, and specifically that in communication networks. However, apart from ecological motivations, there are two additional important drivers for reducing the electricity consumption in backbone networks specifically. First, because most backbone network equipment (such as Internet Protocol (IP) routers) is densely concentrated in telecom operator buildings, which presents major technical challenges with respect to proper heat dissipation, and induces high cooling costs. Second, since the beginning of this millennium the price per unit of electrical energy has started to rise, breaking with the preceding trend where electricity was getting cheaper each year (trend numbers corrected for inflation). This has strong implications for those businesses where electricity consumption is a significant operational cost, such as telecom operators and data center owners.

In this dissertation we contributed to the research on reducing electricity consumption in backbone networks in three ways. First, **we distilled a list of representative power consumption values for backbone equipment, which can be used in current and future research**. This addresses the issue where different studies use(d) widely different values for similar equipment. The underlying data for these representative values have been made publicly available in a report and an online database (http://powerlib.intec.ugent.be). Through the use of an analytical model, we also confirmed that IP routers are indeed the major consumers in current backbone networks, accounting for more than half of the total backbone power consumption.

Second, we evaluated optical circuit switching (i.e., optical bypassing the power-hungry IP routers) as a means to reduce the power consumption. Our results show that circuit switching is always preferable when the average node-to-node demands are higher than half the transport linerates. However, packet switching can become preferable when the traffic demands are lower than half the transport linerates. Apart from the findings on the power saving potential, a key takeaway message is that **the ratio between the average demand and the transport linerate is thus a critical factor to take into account for future, related research**.

Third, we used our analytical model from above to perform a quantitative survey of different power saving approaches for backbone networks. The use of the analytical model allows us to differentiate between and isolate the impact of different techniques, as well as estimate the combined power saving potential. Our results indicate that the power reduction potential of static, once-off approaches ranges from 2.3× (Moderate Effort scenario) to 31× (Best Effort scenario). Factoring in historic and projected yearly efficiency improvements ("Moore's law") roughly doubles both saving potentials on a 10 year horizon. Considering that the traffic is projected to increase close to 10-fold in the next 10 years, **the power reduction potential of our Moderate Effort scenario (6.0×) will not be sufficient to halt the power consumption increase in backbone telecommunication networks**. **However**, the estimated 31–64× reduction for our Best Effort scenario shows that **enough potential is likely available to halt, and even reduce, the absolute power consumption**. This will without doubt require significant efforts from an economic, technical and operational point of view.

Reducing carbon emissions in a distributed data center, using solar and wind energy In the Appendix of this dissertation, we also report related Green ICT work where the focus is shifted from backbone networks to data centers and from electricity consumption to carbon emissions. From an environmental point of view, the drawback of the growing electricity consumption of ICT is mainly linked to the associated carbon emissions. An alternative approach to reducing carbon emissions is to use electricity with a low carbon footprint (such as generated through solar or wind power). In the context of data centers this creates interesting opportunities, as data centers can be located close to those sites which are optimal for renewable power generation. In addition, if we consider a distributed data center consisting of different sites at large geographical distances, computation jobs and data can be shifted to those sites where renewable energy (sun, wind) is available at that point in time. This has been referred to as a Follow The Sun/Follow The Wind (FTSFTW) scenario.

We evaluated such a FTSFTW scenario, taking into account the carbon emissions associated with the manufacturing and operation of the distributed data center. Our contribution is that we show that the manufacturing carbon footprint is a non-negligible factor in this scenario, but—under certain conditions—**minor carbon footprint savings are possible when de-**

ploying *additional* data center sites to fully exploit the geographic availability of renewable energy. However, **larger footprint savings are possible when applying the FTSFTW scenario to a distributed data center where the nominal load is far below the maximum capacity**. We should note that the regional carbon emission intensity of the electricity is a critical factor; countries where the electricity production is not associated with large carbon emission will not benefit from such a FTSFTW scenario.

Future research challenges As with most types of research, this dissertation opened yet another box full of new questions and additional things to study more deeply. We would like to point a few interesting trends and directions below.

Related to our ICT footprint study, it would certainly be interesting to consider to perform a more complete **Life Cycle Analysis (LCA)**, whereby not only the electricity consumption resulting from operation is considered, but also the energy to manufacture, transport, recycle and dispose of ICT equipment. Factoring in the **carbon emissions** associated with each life cycle phase would be the next logical step to take. Furthermore, we only considered the ICT equipment, and not the **ICT services**, such as software development; we can currently only guess that the resulting energy/carbon footprint will be bumped up significantly. Finally, regardless of the above extensions, **periodic estimates** of the worldwide ICT electricity consumption will be essential to provide timely feedback if indeed ICT consumption remains relatively small, or instead continues to grow at an unsustainable rate.

With respect to power consumption in backbone networks, ample interesting research is available. **New equipment** with additional features is becoming available (such as bit rate variable transponders), and the power consumption of network equipment will probably become increasingly proportional to the traffic load; this creates a need to update the set of representative power consumption data with sufficiently accurate power models to capture these developments. More **load proportionality** will very likely also impact the trade-off between circuit switching and packet switching. While we have identified the (rough) power saving potential of individual approaches, many of the approaches have a **cost, implementation and operational aspect** which has not yet been sufficiently evaluated. That would be an important step for taken implementation decisions.

Finally, our evaluation of a Follow The Sun/Follow The Wind scenario for data centers was limited to the associated carbon emissions. However, our model can rather easily be used or modified to evaluate **other metrics**, such as low and high energy prices. With the foreseen (or at least, required) shift to an increasing amount of renewable energy, such evaluations might soon become relevant.

A

Distributed Computing for Footprint Reduction by Exploiting Low-Footprint Energy Availability

The focus of this chapter is slightly different from the previous chapters. We look at data centers instead of backbone networks, and not only consider the electricity consumption but expand it to carbon emissions instead.

⋆ ⋆ ⋆

Ward Van Heddeghem, Willem Vereecken, Didier Colle, Mario Pickavet, and Piet Demeester

Published in Future Generation Computer Systems, February 2012

Abstract Low carbon footprint energy sources such as solar and wind power typically suffer from unpredictable or limited availability. By globally distributing a number of these renewable sources, these effects can largely be compensated for. We look at the feasibility of this approach for powering already distributed data centers in order to operate at a reduced total carbon footprint. From our study we show that carbon footprint reductions are possible, but that these are highly dependent on the approach and parameters involved. Especially the manufacturing footprint and the geographical region are critical parameters to consider. Deploying additional

data centers can help in reducing the total carbon footprint, but substantial reductions can be achieved when data centers with nominal capacity well-below maximum capacity redistribute processing to sites based on renewable energy availability.

A.1 Introduction

Data center power consumption is significant, and growing The last decade has seen a steady rise in data center capacity and associated power consumption. In 2008, the yearly average worldwide data center power consumption was estimated to be around 29 GW [1]. This is comparable to the total electricity consumption of Spain in the same year [2], a country that ranks in the top 15 of the list of electricity consumption per country. In [3], it was estimated that the aggregate electricity use for servers worldwide doubled over the period 2000 to 2005. With the predicted growth of Internet-based services for social networks and video, and with the growing usage of mobile thin clients such as smart phones that require a server back-end [4], it seems unlikely that this increase will halt soon.

Using renewable energy, in addition to energy-efficiency, is key to mitigate climate change While the growing energy consumption in data centers presents some issues both economically and technically, there has been a growing concern from an environmental point of view as well, with electricity consumption contributing to Greenhouse Gas (GHG) emission. Two high-level approaches can help in reducing GHG emissions: (a) an improvement in energy-efficiency to reduce the amount of electrical energy used, and (b) use of energy that contributes little to GHG emissions. What concerns the latter, this electrical energy will typically come from renewable energy sources such as solar and wind power.

Adding renewable energy to the current energy mix still poses some issues While renewable energy is indeed already promoted and used to mitigate climate change both in Information and Communication Technology (ICT) and non-ICT sectors, significantly increasing the amount of renewable energy as part of the regular energy mix raises a number of issues [5]. First, because most good sites for renewable energy sources may be located in distant areas with limited transmission capacity, and it might take many years for the required transmission infrastructure to become available [6]. Second, the distributed power generation poses many challenges for the existing distribution infrastructure, especially with respect to protection and control strategies due to new flow patterns [6] [7]. Third, with renewable energy sources likely to be located in distant areas, the transmission losses will increase; current transmission losses are already estimated to be around 6.5 % of the total electricity disposition [1] for the U.S.A in 2007 [8]. Fourth, with hydro power usually reserved for peak power handling [9], other renewable energy sources such as wind and solar power are usually characterized by intermittent power delivery, resulting in periods of peak power being available and no power being available at all.

Data centers are uniquely positioned to provide an alternative solution Data centers have become more and more globally distributed for a number of reasons as summarized by [10]: "the need for high availability and disaster tolerance, the sheer size of their computational infrastructure and/or the desire to provide uniform access times to the infrastructure from widely distributed client sites". This geographical distribution of data centers, combined with the availability of low-power and high-speed optical links, allows them to be located near renewable energy sites. With technology currently available to migrate live virtual machines while minimizing or avoiding downtime altogether [11] [12] [13], jobs can be dynamically moved from a data center site where renewable power dwindles to a different site with readily available renewable power. This approach has previously been referred to as Follow The Sun/Follow The Wind (FTSFTW) [5].

Fig. A.1 illustrates this concept with solar powered data center sites. As the sun sets in the top-right data center (and the capacity of potential backup-batteries fall below a critical value) the site's data and jobs are moved to a different site (top left) where solar power has become available.

In this paper we will evaluate the carbon footprint and potential footprint savings of such a FTSFTW-based distributed data center. We will generalize on the notion of renewable energy, and instead consider low-footprint (LF) energy and high-footprint (HF) energy. As a metric for the carbon footprint we will use grams of CO_2-eq, unless otherwise indicated. CO_2-eq indicates CO_2-equivalent emissions, which is the amount of CO_2 that would have the same global warming potential when measured over a given time horizon (generally 100 years), as an emitted amount of a long-lived GHG or a mixture of GHG.

The contributions of this paper are the following:

---

[1] To be correct, the losses percentage is calculated as a fraction of the total electricity disposition excluding direct use. Direct use electricity is electricity that is generated at facilities that is not put onto the electricity transmission and distribution grid, and therefore does not contribute to transmission and distribution losses [8].

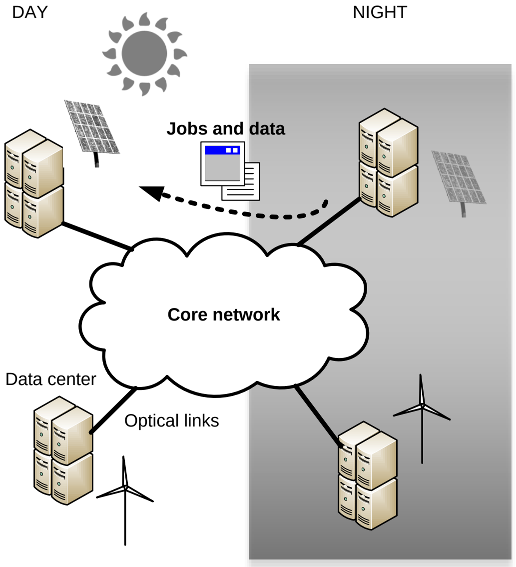

Figure A.1: *Distributed data center*

- we provide a mathematical model for calculating the carbon footprint and savings of such a distributed data center infrastructure which is powered by a fixed mix of LF and HF energy (Section A.3),

- we provide a detailed and realistic quantification of the parameters in our mathematical formulation (Section A.4),

- we show that the manufacturing carbon footprint is a non-negligible factor in footprint reduction evaluations, and that — under certain conditions — minor footprint savings are possible when deploying *additional* sites where jobs are distributed according to the FTSFTW approach (Section A.5),

- we show that larger relative footprint savings are possible when applying the FTSFTW scenario to distributed data centers where the nominal load is well below the maximum capacity (Section A.6).

It should be noted that the theoretical model we present in Section A.3 can be applied, with or without slight modifications, using other metrics than carbon footprint.

A.2 Related work

Next to the work already pointed out in the previous section, below are some earlier references and publication related specifically to the FTSFTW approach.

One of the first papers to suggest locating data centers near renewable energy sources is [14]. The primary reason given is that it is cheaper to transmit data over large distances than to transmit power. The paper does not discuss or explore this issue in any more detail.

The first paper to our knowledge to discuss and mathematically evaluate load distribution across data centers taking into account their energy consumption, energy cost (based on hourly electricity prices) and so-called low-footprint ’green energy’ and high-footprint ’brown energy’ is [10]. It presents and evaluates a framework for optimization-based request distribution, which is solved using heuristic techniques such as simulated annealing. The paper shows that it is possible to exploit green energy to achieve significant reductions in brown energy consumption for small increases in cost. It does not consider the manufacturing carbon footprint.

Similarly, in [15] load distribution across data centers is discussed, but only to optimize energy costs by exploiting energy price differences across regions.

In [5] the FTSFTW scenario is discussed in more detail and an Infrastructure as a Service (IaaS) approach is suggest to turn this in a viable business model. It outlines the main arguments for employing such a scenario. The key idea put forward is that the FTSFTW scenario provides a ’zero-carbon’ infrastructure for ICT, thereby somewhat optimistically ignoring the potential contribution of the manufacturing carbon footprint.

The GreenStar Network project [16] is a proof of concept testbed for the FTSFTW strategy. The project started in 2010 and is deployed across the Canadian-based CANARIE research network and international partners distributed across the world. It consists of a number of small-scale ’nodes’ powered by renewable energy (especially hydro, solar and wind power) which provide energy for the routers, switches and servers located at the node. Applications are running inside virtual machines, with multiple virtual machines per server, and are migrated live from node to node. The expected outcome of the project is a number of tools, protocols and techniques for deploying ’green’ ICT services.

A framework for discovering carbon-minimizing resources in networks similar to those deployed by the GreenStar Network project, is described in [17], but again the manufacturing carbon footprint is not considered.

A.3 Theoretical model

In this section we will outline the details of the scenario that we consider and develop a theoretical model for estimating its total carbon footprint. The quantification of the various parameters in our formulation will be done in Section A.4.

To introduce our theoretical model, we consider the distributed simplified data center infrastructure that is shown in Fig. A.2. It consists of m equally-sized sites. Of these m sites, on average n sites are active. When a specific site becomes non-active, data and processing is moved to another active site, keeping the number of active data centers equal to n at all times.

At this point it is important to point out that, although we use the term data center, our model will be independent of the size of the data center. A data center site could be an energy-optimized building housing thousands of servers, or it could be as small as a single server. In the context of this paper, it might be helpful to think of a data center site as a computing node of any possible size.

Each site is powered by either LF or HF energy. The average availability of LF energy versus HF energy is considered equal, but uncorrelated, for each site. This availability ratio p might be the result of an average temporal availability of a specific renewable energy source (for example, solar or wind power), or specific service level agreements between the data center operator and the utility provider.

To reduce the total footprint, the usage of LF energy will be maximized by migrating operation of a data center powered by HF energy to a data center where LF energy is available. When no LF energy is available, HF energy will be used to guarantee service delivery.

The total carbon footprint F of the above described distributed data center infrastructure, averaged over a long-enough period, will be the sum of the manufacturing footprint F_m, the usage footprint F_u and the communication footprint F_c:

$$F = F_m + F_u + F_c \tag{A.1}$$

The manufacturing footprint will be the carbon emitted during the manufacturing of the sites and the equipment (servers, network equipment etc.) inside. The usage footprint will be the result of the electrical energy used during the use phase. The communication footprint will be the carbon emitted by migrating data and jobs from site to site. All three footprints will be expressed in g CO_2-eq.

Before we elaborate on each of these footprints, it is useful to point out the following assumptions we will make for our theoretical model:

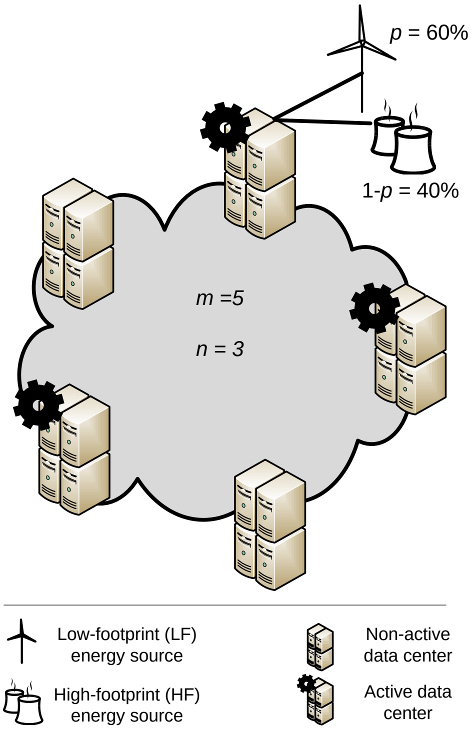

Figure A.2: *Distributed data center infrastructure overview, consisting of* m*=5 sites with* n*=3 sites active. The independent LF energy availability per site is* p*=0.6*

- We assume each site in the distributed data center to be of uniform size.
- We assume instant site migration. That is, we assume that a migration takes no time and produces no extra overhead not accounted for in the communication footprint. If the migration frequency is relatively low (say, limited to a few times a day), this assumption will hold.
- We do not consider a surplus of LF energy. That is, if for example 4 out of 5 sites have LF renewable energy available, but we only require 3 sites for daily operation, the electricity generated in the 4th site is 'wasted'. There is potential for using this energy for other less-critical purposes, or for selling or trading it for carbon credits. However, for simplicity and generality, our model does not take using surplus available power into account.
- We assume that a non-active data center site consumes no energy. While this is an optimistic assumption for large data centers, this is certainly feasible for micro-scale data centers consisting of a few servers (remember that, although we use the term data center, our model is independent of the data center size). The energy for a non-active site could be reduced to (nearly) zero by for example suspending all servers.

A.3.1 Usage footprint

Let's call p the chance that a site is powered by LF energy. Let's call k the total number of data center sites that are powered by LF energy and P_k the chance of this number being k. This chance is given by the probability mass function of the binomial distribution:

$$P_k = \binom{m}{k} p^k \left(1 - p\right)^{m-k} \tag{A.2}$$

Eq. (A.2) can be understood intuitively as follows. The chance for exactly k sites powered by LF energy is p^k. The chance for the $m-k$ remaining sites to be *not* powered by LF energy is $\left(1-p\right)^{m-k}$. The number of ways to choose k sites out of a total of m sites is given by the binomial coefficient $\binom{m}{k}$ and can be calculated as $\binom{m}{k} = \frac{m!}{k!(m-k)!}$.

Given L the carbon footprint of the total usage phase of a single site when powered exclusively by LF energy and H the carbon footprint when powered exclusively by HF energy. The total usage footprint F_u for all sites is then:

If $k \geq n$ (that is, if LF energy is available in enough or more sites than required):

$$F_u = nL \tag{A.3}$$

Else:

$$F_u = (n-k)H + kL \tag{A.4}$$

Thus, using the chances of k being a certain value, the total usage footprint F_u becomes:

$$F_u = \sum_{k=n}^{m} [P_k nL] + \sum_{k=0}^{n-1} [P_k ((n-k)H + kL)] \tag{A.5}$$

The first term describes the weighted footprint if enough sites are powered by LF energy, the second term when this is not the case. When substituting Eq. (A.2) in Eq. (A.5) we get for the total usage footprint F_u of the distributed data center infrastructure:

$$\begin{aligned} F_u &= nL \sum_{k=n}^{m} \left[\binom{m}{k} p^k (1-p)^{m-k} \right] \\ &\quad + \sum_{k=0}^{n-1} \left[\binom{m}{k} p^k (1-p)^{m-k} ((n-k)H + kL) \right] \end{aligned} \tag{A.6}$$

The usage footprint results exclusively from electrical energy. The emission intensity of electricity describes the GHG emissions in gram CO_2-eq per kWh. We use I_L and I_H to denote the emission intensity for LF and HF electricity respectively. With E_u the energy used by a single site during the entire use phase, L and H can thus be expressed as:

$$\begin{cases} L = I_L E_u \\ H = I_H E_u \end{cases} \tag{A.7}$$

A.3.2 Manufacturing footprint

The total manufacturing footprint F_m is a function of the carbon footprint cost M for manufacturing one data center site, and the number of data centers sites m:

$$F_m = mM \tag{A.8}$$

Table A.1: *Manufacturing fraction values according to different studies*

Reference	Description	Manufacturing phase	Use phase (4 years)	f
PE Int. [18]	Office server [a]	500 kg CO_2-eq/unit	1030 kg CO_2-eq	0.49
Malm. [19]	PC [a]	400 kg CO_2-eq/unit	640 kg CO_2-eq	0.63
Malm. [19]	Server	500 kg CO_2-eq/unit	5200 kg CO_2-eq	0.10
Malm. [20]	Data centers[b]	10 Mton CO_2-eq in 2007	108 Mton CO_2-eq in 2007	0.09

[a] Overhead power in use phase not included (PUE=1). See text for more information

[b] This includes data center equipment and buildings. Data based on 10 million new servers and 35 million servers in use; this translates roughly to a use phase of 4 years. Use phase emission intensity in [20] = 0.6 kg CO_2-eq/kWh

As we will see in Section A.3.4, it is convenient to consider the manufacturing fraction f, which is the ratio of the manufacturing carbon footprint M of a single site over the usage carbon footprint H of a single site:

$$f = \frac{M}{H} \tag{A.9}$$

Equipment where the manufacturing emits less GHG than the typical GHG emitted during its use phase will have a manufacturing fraction $f < 1$.

Given Eq. (A.9), we can rewrite Eq. (A.8) as:

$$\begin{aligned} F_m &= mfH \\ &= mfI_H E_u \end{aligned} \tag{A.10}$$

Note that we considered the equipment to be manufactured with HF energy, by expressing M as a function of H instead of L.

A.3.3 Communication footprint

Migrating jobs or data across data centers incurs an extra amount of carbon emissions. This will mainly be due to the energy consumed for (a) the transportation over an optical network, (b) the preparation and duration of the migration and (c) switching the data center to the non-active state or vice versa. In this section we show that the overhead of the above three factors is negligible with respect to the carbon emitted in the manufacturing and use phase, and can thus be ignored for now.

Data centers are typically connected by optical networks. Power consumption in the optical core network is dominated by the IP router power consumption, with high-end IP routers consuming in the order of 10 W/Gbps [21]. Accounting for redundancy, cooling and power supply overhead, and client and network interface, we have approximately 100 W/Gbps, or an energy of 2.7 10^{-5}kWh needed to transport one Gbit.

Further, we assume two migrations per site once a day, i.e. one inbound migration and one outbound migration. We consider each server in a data center site to be capable of running four virtual machines, with each virtual machine to be about 10 Gbyte in size. For each server's data to be migrated, this totals to 640 Gbit/day. Considering a server use phase of 4 years, this sums up to 934 000 Gbit per use phase. Using our estimation from above, this requires approximately 26 kWh of energy. With a world-average emission intensity of 500 g CO_2-eq/kWh, this results in about 13 kg CO_2-eq emitted due to migration (for one server, during its entire use phase). This equals to less than 3% of the current manufacturing footprint of a server (about 500 kg CO_2-eq, see Table A.1), or about 0.5% of the current total carbon emissions.

With respect to the energy overhead induced by migration preparation and duration, transmitting our exemplary 640 Gbit/day would take less than 15 minutes per day over a 1 Gbps link. This accounts for only about 1% of the time.

Likewise, as the daily migration frequency is low, the time and energy overhead to switch a data center from the active to non-active state (or vice versa) should be relatively low as well. Also, the active/non-active switchover time will probably depend on the kind of jobs and data that the data center is running.

Although the above estimate is based on the current situation of the average absolute carbon footprint of servers and current virtualization technology, we feel that it is a fair assumption for current and short term future to neglect the contribution of the communication footprint F_c to the total footprint.

A.3.4 Total footprint

Combining Eq. (A.6) and Eq. (A.10), the total footprint is given by:

$$\begin{aligned} F &= mfI_HE_u \\ &\quad +nL\sum_{k=n}^{m}\left[\binom{m}{k}p^k\left(1-p\right)^{m-k}\right] \\ &\quad +\sum_{k=0}^{n-1}\left[\binom{m}{k}p^k\left(1-p\right)^{m-k}\left((n-k)H+kL\right)\right] \end{aligned} \tag{A.11}$$

The above equation depends on the value of E_u, the single site usage energy. This value will vary depending on the data center size and type, and on the jobs and data processed. We can eliminate this parameter, if we normalize the total footprint over the single site usage energy E_u.

By doing so, we can conveniently express this total normalized footprint F_{norm} as a function of the LF energy emission intensity I_L, the HF energy emission intensity I_H and the fraction f:

$$\begin{aligned} F_{norm} &= \frac{F}{E_u} \\ &= mfI_H \\ &\quad +nI_L \sum_{k=n}^{m} \left[\binom{m}{k} p^k (1-p)^{m-k} \right] \\ &\quad + \sum_{k=0}^{n-1} \left[\binom{m}{k} p^k (1-p)^{m-k} \left((n-k)I_H + kI_L \right) \right] \end{aligned} \tag{A.12}$$

We now have a metric for the carbon footprint which is independent from the data center size and type, and with unit [g CO_2-eq/kWh].

A.4 Parameter quantification

Our model constructed in the section above consists of a number of parameters. In this section we discuss realistic values for each of these parameters.

A.4.1 Manufacturing fraction (f)

The manufacturing fraction represents the ratio between the manufacturing carbon footprint and the usage carbon footprint. Detailed Life Cycle Analysis (LCA) studies that report on the carbon emissions of data centers during the manufacturing phase and the use phase are scarce. Moreover, the resulting manufacturing fraction is influenced by the use phase lifetime of the equipment and the emission intensity of the energy used during the use phase. In addition, it is important to know if reported use phase values include power consumed for overhead such as cooling. This overhead is typically expressed by the Power Usage Effectiveness (PUE). For example, a PUE of 2 (a typical accepted value for data centers[2]) indicates that for each Watt consumed by useful equipment such as servers and switches an additional Watt is consumed through overhead.

---

[2] Recently deployed high-capacity data centers with a focus on energy efficiency show much lower PUE values, such as Google claiming to reach a yearly average of 1.16 at the end of 2010 [22]. However, as the LCA data is based on 2007 estimates, the for that year typically accepted PUE value of 2 is used [23].

Table A.1 lists emission values and the derived manufacturing fraction f according to a number of studies. All data, except for the 'Simple office server' and the 'PC', includes overhead power consumption. For the 'Simple office server' probably no overhead is included ([18] isn't completely clear on this); correcting for this with a PUE of 2, the use phase power consumption doubles and thus the manufacturing fraction value halves, bringing the values roughly in line with the other data.

Based on the data in Table A.1 we will use, unless otherwise specified, a value of f=0.25.

A.4.2 High-footprint energy emission intensity (I_H)

The parameter I_H indicates the emission intensity of regular (HF) electrical energy. As already stated, the emission intensity indicates the amount of GHG emitted for each kWh of electrical energy, and is typically expressed in grams of CO_2-eq per kWh.

The value for I_H differs from country to country, and for larger countries even from region to region, depending on the primary energy sources (such as coal or gas) and technologies (such as open cycle gas turbines or combined cycle gas turbines) used for generating electricity, see for example Table A.2 [3].

For this paper, we will consider the world average value of 500 g CO_2-eq/kWh.

Table A.2: *Average CO_2 emissions per kWh from electricity and heat generation for a number of countries and regions, data for 2008 [24]*

Region	Intensity [g CO_2/kWh]
World	502
United States	535
Canada	181
European Union	351
China	745
India	968

A.4.3 Low-footprint energy emission intensity (I_L)

The emission intensity I_L for low-footprint electricity is obviously lower than the regular HF energy emission intensity I_H. Indicative, Fig. A.3 lists

---

[3] The table reports the CO_2 emissions instead of the CO_2-eq emission (which takes a number of other GHG into account). However differences are minor and irrelevant for our study

the estimated emission intensity for a number of low-footprint sources (typically renewable energy such as hydro, wind or solar power), as reported by [9]. Roughly similar numbers are given in the slightly older study of [25].

In this paper, we assume a state-of-the-art LF energy emission intensity of 10 g CO_2-eq/kWh.

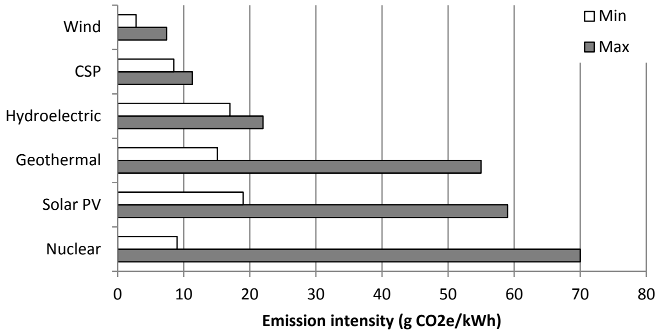

Figure A.3: *Lower and upper emission intensity estimates for various low-footprint sources [9] (CSP: Concentrated Solar Power)*

A.4.4 Low-footprint energy availability (p)

The parameter p represents the chance of each site being powered by LF energy. For example, with p=0.6, each site has an independent chance of 60% to be powered by LF energy at any point in time. Or otherwise put, 60% of the time, each site will be powered by LF energy.

While it might seem tempting to try to relate the value for p to the availability of a specific LF energy source (say, wind energy), this is not necessary for our model. After all, the availability of LF energy sufficient for powering a data center site will largely be a matter of monetary cost. This cost will be reflected either in the negotiated service level agreement (SLA) with the utility provider, or in the cost to install the required capacity of LF energy sources to deliver the required nominal power even during periods of low availability of e.g. sun or wind. Thus, a higher value for p will usually require higher investments. Note that it is key for the validity of our footprint model to known what kind of power (LF or HF) is used at what point in time, so as to be able to migrate the data to a different site if needed (and if possible).

We assume p=0.6, as we will see later that this results in maximum savings.

A.5 Case study I: The Added Distributed Data Centers (ADD) Scenario

Can we reduce the footprint of a regular data center, by distributing additional sites across the globe as to benefit from uncorrelated and potentially complementary availability of renewable energy sources which offer a lower usage footprint? This is the question we will examine in this section. We refer to this scenario as the *Added Distributed Data centers (ADD) scenario*.

Consider a data center that requires n=3 sites for daily operation. Each site has an LF energy availability of p=0.6, and we consider the current estimation for the manufacturing fraction f=0.25. Since we want to reduce the footprint of the complete data center, we would like to be able to run our applications on three data centers that have LF energy available. The chance of success increases with an increased number of data centers to choose from, that is, if we increase the total number of sites m to a value higher than 3.

Fig. A.4 shows the use phase, manufacturing phase and total footprint as we increase the total number of data centers m beyond 3. With each additional data center, the use phase footprint decreases as a result of the increased chance of finding a data center that runs on LF energy. Initially, this decrease is large enough to make up for the linearly increasing manufacturing footprint, resulting in a decreasing total footprint. However, when the number of data centers is approximately the double of the number of data centers required, the total footprint increases and eventually overtakes the first scenario footprint.

Taking the first scenario (where m=n=3) as a baseline, we see initial footprint savings until too much data centers are deployed, resulting in a net loss. Taking the first scenario as the baseline makes sense, since this corresponds to the current practice of operating a number of sites with a mix according to p of LF and HF energy, *without* migrating data or processing capacity based on LF energy availability.

A.5.1 Influence of manufacturing fraction (f)

As we have seen in the above case, the usage footprint reduction was initially able to make up for the linearly increasing manufacturing footprint. What if the manufacturing fraction f is higher, say f=0.5? Fig. A.5 shows the normalized footprint (upper figure) and relative savings (lower figure) for different values of f.

Clearly, footprint reduction becomes smaller and even impossible for

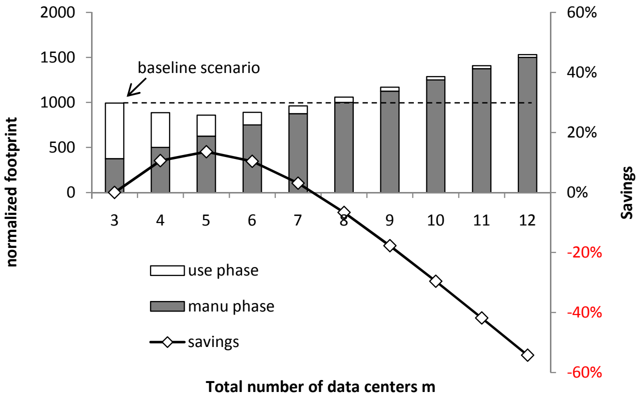

Figure A.4: *The total normalized footprint F_{norm} and corresponding relative emission savings as a function of the total number of data centers m. Savings are calculated with respect to the baseline scenario. (Parameter values: n=3, f=0.25, p=0.6, I_L=10 g CO_2-eq/kWh, I_H=500 g CO_2-eq/kWh and E_u=1)*

higher values of f. Even more so, our current rough estimate of f=0.25 seems critical: with a slightly higher value for f=0.3 savings are almost negligible (a mere optimistic 5%) and might be completely annihilated if we take more subtle factors (such as the migration footprint and management overhead) into account.

In the inverse case, for lower values of f the savings increase. At the utopian case of having manufacturing for free (f=M=0), savings are obviously maximal and converge to the usage footprint cost nL.

A.5.2 Influence of low-footprint energy availability (p)

Perhaps counterintuitive, an increase of LF energy availability of p towards 100% does not unconditionally result in additional savings. While the footprint indeed decreases monotonic with an increase of p (because the usage footprint becomes smaller), the baseline scenario footprint (where m=n) will also decrease.

Fig. A.6 shows that for the scenario n=3, m=6 (i.e., twice as much data centers as required for daily operation) the savings are maximum around p=0.5 to 0.6. For p=0 there is a net loss due to the increased manufacturing footprint not yet being offset by a greener usage footprint. For p=1 the baseline scenario runs entirely on LF energy whereas the FTSFTW approach has an increased manufacturing footprint due to the extra sites deployed.

As we have already argued that p will be cost driven, a case-based cost

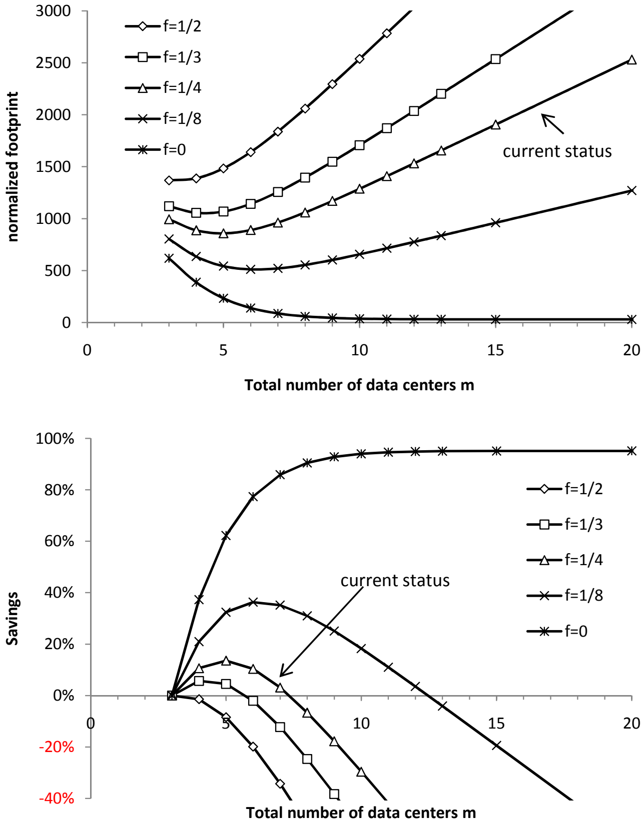

Figure A.5: *The total normalized footprint F_{norm} and relative emission savings for n=3 as a function of m for different manufacturing fractions f*

study will have to find the optimal value for p. In retrospect, this also explains our decision for taking p=0.6.

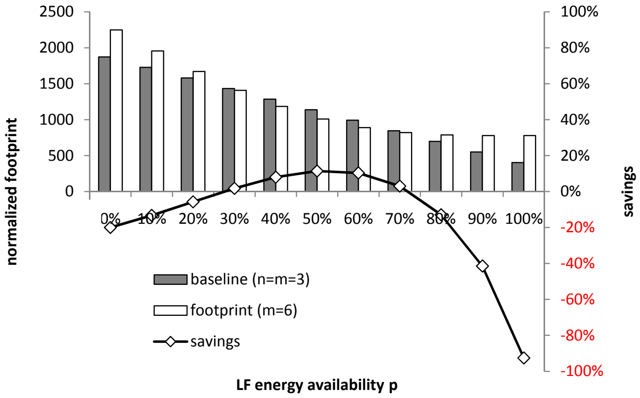

Figure A.6: *The total normalized footprint and the relative savings (with respect to the baseline scenario where m=n=3) as a function of the LF energy availability p (n=3, m=6 and f=0.25)*

A.5.3 Influence of n and m values

Because of the binomial coefficient, we cannot simply generalize the footprint savings obtained for e.g. n=3 and m=6 to apply to any other combination of n and m with the same ratio, e.g. n=1 and m=2, or n=10 and m=20.

For higher values of n, footprint savings already occur for higher (i.e., worse) manufacturing fractions. For example, when we consider n=10 (see Fig. A.7), already for f=0.5 minor savings are available (2% maximum), whereas for the previous case where n=3 this was not the case (see Fig. A.5). Because of the higher number of sites, the chance for finding enough sites where LF energy is available has increased. It should be noted that the *total* footprint will have increased as well.

This finding suggests to favor a large number of small, distributed data center sites, over a few large ones. However, in that case, care should be taken that the combined manufacturing footprint of the small sites is not larger than the manufacturing footprint of the few larger sites. Taking the idea to extremes, large-scale distributed computing projects such as Folding@Home [26] where small consumer entertainment devices are involved [27] might be a perfect fit, if both the manufacturing footprint and usage footprint (standby power consumption issues etc.) from these devices is small enough.

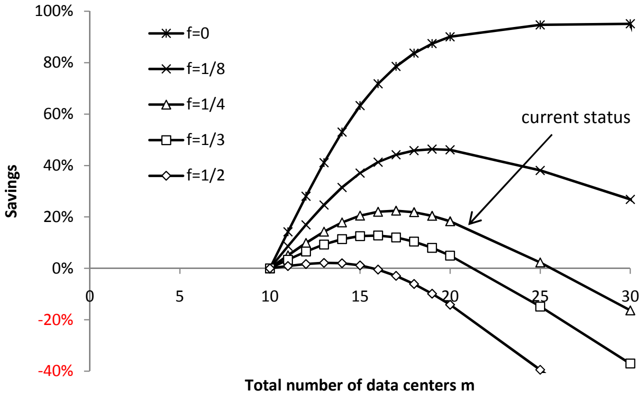

Figure A.7: *Relative emission savings for n=10 (all other parameters are equal as before)*

A.5.4 Influence of emission intensity difference

The HF emission intensity (500 g CO_2-eq/kWh) en LF emission intensity (10 g CO_2-eq/kWh) that we consider in this paper following our findings in Section A.4.2 and Section A.4.3 are relatively large in difference; I_L is only 2% of the I_H. In some countries or regions, the regular emission intensity is substantially lower (or higher) than the world average value, as can be seen in Table A.2. Will the FTSFTW approach still be sustainable under those conditions?

Fig. A.8 shows the relative savings with changing values of I_H. It is immediately clear from this figure that for values below the world average, the savings quickly become negligible. For emission intensities below the average European value, savings become negative, i.e. *more* carbon dioxide will be emitted. On the contrary, for geographical regions where the regular electricity has high emission intensities (such as China and India), the savings offered by FTSFTW are much higher.

Note that we consider the manufacturing carbon footprint cost M (see Eq. (A.8)) to be fixed, even with changing I_H value. This means that in this case we have fixed the instance of I_H in Eq. (A.10) to the world-average emission intensity. Fixing the manufacturing footprint makes sense, as it represents the case where the equipment remains manufactured as before, but is used in a region with a different HF energy emission intensity.

Similarly, we can also consider different values for the LF energy emission intensity I_L. The value of I_L=10 g CO_2-eq/kWh we assumed in Sec-

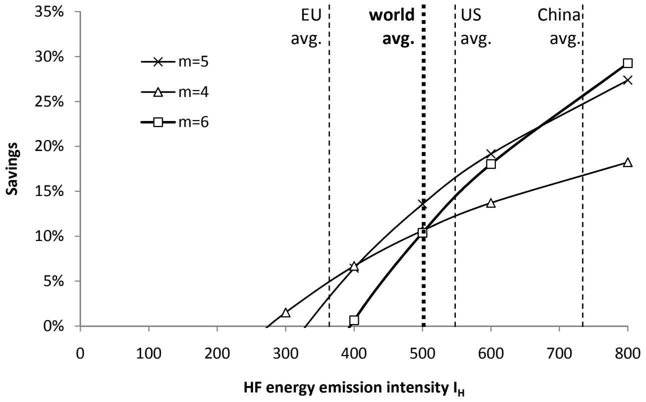

Figure A.8: *Relative emission savings as a function of the HF energy emission intensity I_H. (for f=0.25, p=0.6 and I_L=10 g CO_2-eq/kWh)*

tion A.4.3, is based on state-of-the art renewable energy, typically from wind turbines. For other energy sources with higher emission intensities, the savings will obviously be smaller.

Fig. A.9 shows the savings for increasing values of I_L, with the HF energy emission intensity fixed at 500 g CO_2-eq/kWh. As can be seen, the savings rapidly dwindle, to the point where they become marginal. As such, using less emission-saving renewable energy sources such as solar PV installation should be evaluated carefully if the main goal is saving on total carbon emission by employing the ADD scenario.

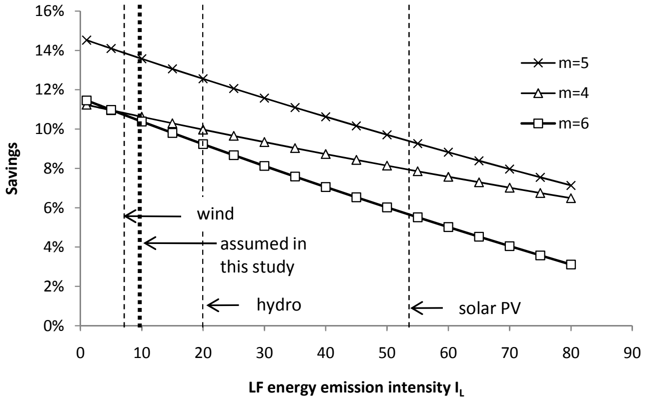

Figure A.9: *Relative emission savings as a function of the LF energy emission intensity I_L. (for f=0.25, p=0.6 and I_H=500 g CO_2-eq/kWh)*

To summarize, with current estimates for the manufacturing and usage footprint, carbon emission savings up to around 14% are possible by deploying additional data center sites. Actual savings depend mainly on the manufacturing fraction (lower is better), the LF energy availability (optimum around 50–70%) and the number of sites deployed (optimum around 1.5 to 2.5 times as much data centers as required for daily operation). For geographical regions with higher HF emission intensities, the possible savings by employing the ADD scenario are much higher than 14%; likewise, for intensities below the world average savings quickly turn negative.

A.6 Case study II: The Low Load Redistribution (LLR) Scenario

The main conclusion from the above scenario is that the manufacturing carbon footprint is a non-negligible factor, and should be taken into account when evaluating potential carbon footprint savings. However, there are cases where the manufacturing footprint is already expended. Data centers are not constantly running at peak capacity, but instead operate at a nominal load well below the peak capacity, typically servers operate most of the time between 10 and 50 percent of their maximum utilization levels [28]. We could redistribute the load using the FTSFTW approach, resulting in what we will refer to as the *Low Load Redistribution (LLR) scenario*.

Regular approach Fig. A.10a shows the regular approach (without applying LLR). The load is equally distributed among the different sites. To calculate the total carbon footprint, we consider a data center with peak capacity m to run at nominal load n. We assume unused servers to be powered down. The total footprint of a data center running at this nominal load is then:

$$F_{nominal} = F_u + F_m = n\,(pL + (1-p)\,H) + mM \tag{A.13}$$

LLR approach What would happen if we apply the FTSFTW approach to optimally distribute processing to sites where LF energy is available (Fig. A.10b)? We can use Eq. (A.11) or Eq. (A.12) to calculate the footprint in that case as well, with m representing the peak capacity, and n representing the (varying) nominal load.

Fig. A.11 plots the footprint for both scenarios for a distributed data center consisting of 5 sites (m=5), for an increasing load (i.e., n increasing from 0 to m). The LF energy availability p per site has been taken equal to

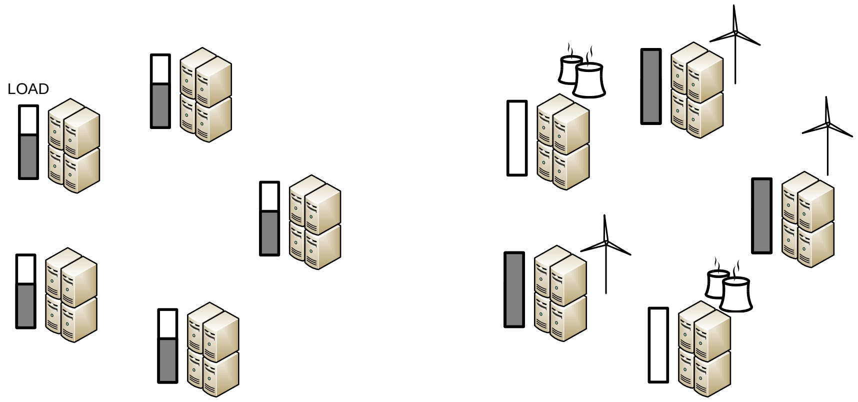

(a) Regular scenario: 3/5 load distribution under the regular approach

(b) LLR scenario: 3/5 load distribution under the FTSFTW approach

Figure A.10: *Nominal load distribution in a distributed data center. (a) shows the regular scenario where a nominal load of 60% is distributed equally over all data center sites. (b) shows the LLR scenario, where the same nominal load is distributed according to the FTSFTW approach, resulting in an optimal usage of sites with LF energy availability*

0.6. As can be seen, for the nominal load being half of the peak capacity, savings around 20% are possible by employing FTSFTW. These are savings over the total footprint, that is, the sum of the use phase and manufacturing phase footprint.

If we only consider the savings over the usage phase, which would be an equally valid approach since the manufacturing phase has no savings, the savings are as high as 90% when running at 20% of the capacity and still reach more than 60% when running at half the peak capacity.

The savings itself vary for different values of p. This is shown in Fig. A.12.

It is important to remark that from the above results we should not conclude to design distributed data centers to run well below their maximum capacity. This results in a large total manufacturing carbon footprint. First, and foremost, data center capacity should be scaled to their nominal loads as much as possible, taking into account such factors as redundancy and peak loads. Once this is done, carbon emissions can be reduced using the LLR scenario outlined above.

A.7 Conclusions

The carbon footprint from data centers is significant, and growing. Besides improvements in energy-efficiency, the use of low footprint energy (typi-

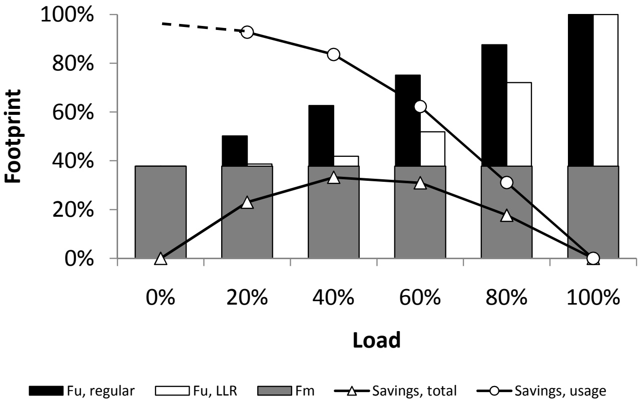

Figure A.11: *Relative footprint (with respect to the maximum load) of a distributed data center running at various loads both under a regular scenario and a LLR scenario (m=5, p=0.6). The 'Savings, total' are the relative savings over the total footprint (both manufacturing F_m and usage F_u). The 'Savings, usage' are the relative savings over the usage footprint only.*

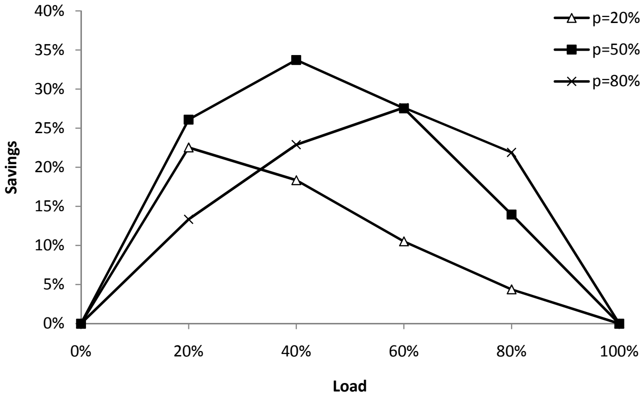

Figure A.12: *Savings for various values of p* cally from renewable energy sources such as wind or solar power) is key to reducing data center carbon emissions. Data centers are in a unique position to overcome some of the issues currently associated with renewable energy sources. They can be located near renewable energy sites, and jobs and data can be migrated from site to site as renewable energy — intermittent by nature — comes and goes . This approach has been referred to as follow the sun/follow the wind.

In this paper, we researched if carbon emissions can be reduced by applying this technique to take advantage of the resulting increased availability of low footprint renewable energy. To this purpose, we have build a mathematical model to calculate the carbon footprint of such a distributed data center infrastructure that is powered by a mix of low-footprint (LF) and high-footprint (HF) energy. We have shown that for footprint reduction the manufacturing carbon footprint of data centers is a critical parameter to consider. Based on the available LCA data for data centers, footprint savings in the order of 14% over the total footprint are possible by deploying additional data center sites to take advantage of the resulting increased available of LF energy. Reductions of the manufacturing footprint relative to the usage footprint will lead to improved savings. However, a number of factors heavily influence the actual savings, which could easily turn into an *increased* carbon footprint if not evaluated carefully. As the savings are strongly influenced by the HF electrical emission intensity, it is of no use to deploy the follow the sun/follow the wind approach in regions with emission intensities below the current world average value. And, consequently, it makes more sense to use the approach in regions with high emission intensities. Carbon footprint savings also depend on the LF energy availability per site: optimal availability varies for different configurations, but is in the order of 50–70%. Optimum savings can be gained at architectures that deploy around 1.5 to 2.5 times as much data centers as required for daily operation.

Bigger savings — up to 60–90% — are possible by applying the follow the sun/follow the wind strategy to data centers where the nominal load is well-below the peak capacity.

Finally, it should be noted that our model is not restricted to carbon footprint metrics. It can easily be used or modified to evaluate other metrics. For example, the low and high emission intensities can be replaced by low and high energy prices (requiring an appropriate quantification of the manufacturing fraction in that case) to evaluate the cost benefits in the light of fluctuating energy prices. However, this is outside the scope of this paper.

Acknowledgments

The authors would like to thank Koen Casier, Kevin Mets and Tom Verschueren for their helpful advice and discussions. We also thank the Greenstar Network partners for their valuable comments. The work described in this paper was carried out with the support of the BONE project (Building the Future Optical Network in Europe) and the TREND project (Towards Real Energy-efficient Network Design), both a Network of Excellence funded by the European Community's Seventh Framework; the IBBT-project GreenICT and the STRONGEST project funded by the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement n. 247674.

References

[1] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt, and P. Demeester. *Worldwide energy needs for ICT: the rise of power-aware networking*. In 2008 2nd International Symposium on Advanced Networks and Telecommunication Systems, page 3 pp., 2008.

[2] EIA. *International Energy Statistics, total electricity net consumption*. http://tonto.eia.doe.gov/cfapps/ipdbproject/IEDIndex3.cfm?tid=2&pid=2&aid=2, 2010.

[3] J. Koomey. *Estimating total power consumption by servers in the US and the world*, 2007.

[4] L. Deboosere, P. Simoens, J. De Wachter, B. Vankeirsbilck, F. De Turck, B. Dhoedt, and P. Demeester. *Grid design for mobile thin client computing*. Future Generation Computer Systems, 2010.

[5] S. Figuerola, M. Lemay, V. Reijs, M. Savoie, and B. St Arnaud. *Converged Optical Network Infrastructures in Support of Future Internet and Grid Services Using IaaS to Reduce GHG Emissions*. Lightwave Technology, Journal of, 27(12):1941–1946, 2009.

[6] A. Ipakchi and F. Albuyeh. *Grid of the future*. Power and Energy Magazine, IEEE, 7(2):52–62, 2009.

[7] T. Ackermann, G. Andersson, and L. Söder. *Distributed generation: a definition*. Electric Power Systems Research, 57(3):195–204, 2001.

[8] EIA. *Total annual losses estimates related to electrical transmission and distribution*. http://tonto.eia.doe.gov/ask/electricity_faqs.asp#electric$_$rates2, 2007.

[9] M. Jacobson. *Review of solutions to global warming, air pollution, and energy security*. Energy & Environmental Science, 2(2):148–173, 2009.

[10] K. Le, R. Bianchini, M. Martonosi, and T. Nguyen. *Cost-and energy-aware load distribution across data centers*. Proceedings of HotPower, 2009.

[11] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield. *Live migration of virtual machines*. In NSDI'05: Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation, pages 273–286, Berkeley, CA, USA, 2005. USENIX Association.

[12] F. Travostino, P. Daspit, L. Gommans, C. Jog, C. De Laat, J. Mambretti, I. Monga, B. Van Oudenaarde, and S. Raghunath. *Seamless live migration of virtual machines over the MAN/WAN*. Future Generation Computer Systems, 22(8):901–907, 2006.

[13] G. Kecskemeti, G. Terstyanszky, P. Kacsuk, and Z. Neméth. *An approach for virtual appliance distribution for service deployment*. Future Generation Computer Systems, 27(3):280–289, 2011.

[14] A. Hopper and A. Rice. *Computing for the Future of the Planet*. Phil. Trans. R. Soc. A, October 2008. Available from: http://www.cl.cam.ac.uk/research/dtg/www/files/publications/public/acr31/hopper-rs.pdf.

[15] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. *Cutting the electric bill for internet-scale systems*. In Proceedings of the ACM SIGCOMM 2009 conference on Data communication, pages 123–134. ACM, 2009.

[16] *The Greenstar Network Project*. Project website: http://www.greenstarnetwork.com/, 2010.

[17] A. Daouadji, K.-K. Nguyen, m. Lemay, and M. Cheriet. *Ontology-based Resource Description and Discovery Framework For Low Carbon Grid Networks*. In SmartGridComm 2010, 2010.

[18] C. Hermann. *Environmental footprint of ICT equipment in manufacture, use and end of life*. In ECOC 2008, Brussels, page 24, 2008. Available from: http://www.ecoc2008.org/documents/ECOCSymposiumPresentationCHerrmann.pdf.

[19] J. Malmodin. *Global Carbon Footprint of ICT: Method and Results*. ITU FG ICT & CC, Geneva, November 2008.

[20] J. Malmodin, A. Moberg, D. Lundén, G. Finnveden, and N. Lövehagen. *Greenhouse Gas Emissions and Operational Electricity Use in the ICT and Entertainment & Media Sectors*. Journal of Industrial Ecology, October 2010.

[21] W. Van Heddeghem, M. De Groote, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester. *Energy-efficiency in telecommunications networks : link-by-link versus end-to-end grooming*. In Conference on Optical Network Design and Modeling, 14th, Proceedings, page 6. IEEE, 2010.

[22] Google. *Data center efficiency measurements*. http://www.google.com/corporate/datacenter/efficiency-measurements.html, 2011.

[23] U.S. Environmental Protection Agency. *Report to Congress on Server and Data Center Energy Efficiency*. http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf, 2007.

[24] IEA. *CO2 Emissions from Fuel Combustion - highlights*. http://www.iea.org/co2highlights/, 2010.

[25] M. Pehnt. *Dynamic life cycle assessment (LCA) of renewable energy technologies*. Renewable Energy, 31(1):55–71, 2006.

[26] Stanford University. *Folding@Home*. Project page: http://folding.stanford.edu/, 2000.

[27] E. Luttmann, D. Ensign, V. Vaidyanathan, M. Houston, N. Rimon, J. Øland, G. Jayachandran, M. Friedrichs, and V. Pande. *Accelerating molecular dynamic simulation on the cell processor and Playstation 3*. Journal of computational chemistry, 30(2):268–274, 2009.

[28] L. Barroso and U. Holzle. *The case for energy-proportional computing*. Computer, 40(12):33–37, 2007.