

Energy-Aware Adaptive Routing Solutions in IP-over-WDM Networks

vorgelegt von
Master of Science
Filip Idzikowski

von der Fakultät IV - Elektrotechnik und Informatik
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Ingenieurwissenschaften
– Dr.-Ing. –

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr.-Ing. Hans-Joachim Grallert
Gutachter: Prof. Dr.-Ing. Adam Wolisz
Gutachter: Prof. Dr.-Ing. Andreas Kirstädter
Gutachter: Prof. dr hab. inż. Wojciech Kabaciński

Tag der wissenschaftlichen Aussprache: 30. Juni 2014

Berlin 2014

D83

Acknowledgements

I explicitly use the personal pronoun “we” instead of “I” in this thesis in order to acknowledge all the people that I collaborated with and everybody who helped me throughout the work on this thesis.

I express my gratitude to Prof. Adam Wolisz for hosting me at Telecommunications Networks Group (TKN) towards the successful PhD defense. My thesis would not be “round” without you, Prof. Wolisz! Thank you for your constructive comments. I am particularly grateful for initiating collaboration with Zuse Institut Berlin, which resulted in the ONDM2010 and the OSN2011 papers. I thank very much Prof. Andreas Kirstädter and Prof. Wojciech Kabaciński, who evaluated my thesis at a high pace. Thank you for your comments which helped me improve the thesis.

I discussed many, many issues about optical networking with Ahmad Rostami and Hagen Woerner. The topics often went beyond optical networks. Thanks a lot for your support, dear Ahmad and Hagen!

The milestone papers (ONDM2010 and OSN2011 papers) would not be possible without Sebastian Orlowski and Christian Raack! I learnt a lot from you guys! I will always be wondering how it was possible to do the nice piece of research without being supported by the same grant, project, action, etc.

The time that I spent at TKN would not be the same without all the TKN-members. Berthold Rathke, Tobias Poschwatta and Marc Emmelmann provided inestimable support in the beginning of my stay in Berlin. Matthias Bohge and Daniel Willkomm provided motivation for thesis writing by organizing “TKN Graduiertenkolleg” and setting their examples. Sven Wiethölter, Jan-Hinrich Hauer and Osama Khader wrote their theses more or less in the same time period as I did. Thank you very much for the collaborative atmosphere and your hints!

I shall never forget Łukasz Budzisz and Irina Piens and their tremendous support to my research. The numerous discussions tackled not only the research, but also the formalities of projects that I was involved in. Thank you very, very much!

I thank also Murad Abusbaih, Konstantin Miller and Thomas Menzel for numerous discussions, friendly atmosphere and help in the organization of the PhD defense.

I would like to acknowledge the students that I worked with at TKN, particularly Mickaël Guth and Francesco Portoso. I highly value your sincere engagement! The technical and formal issues were easier to solve thanks to Petra Hutt, Heike Klemz, Sonja Cekar, Angelika Kaminska, Sven Spuida, Jürgen Malinowski, Peter-René Shröter and Georgios Ainatzes. Thank you very much!

I had the pleasure to collaborate with several international research groups throughout the years that I spent in Berlin. First, I would like to mention Fabio Neri, Marco Ajmone Marsan and Michela Meo for proposing and leading the TREND Network of Excellence (supported by the European Community’s Seventh Framework Programme FP7/2007-2013 under grant agreement n. 257740). Second, I want to explicitly thank the co-authors of the papers that contain main parts of this thesis, i.e., Edoardo Bonetto, Luca Chiaraviglio, Antonio Cianfrani, Angelo Coiro, Raul Duque, Felipe Jiménez, Esther Le Rouzic, Francesco Matera, Francesco Musumeci, Edion Tego,

Ward Van Heddeghem, Jorge López Vizcaíno, and Yabin Ye. Thank you very much for the fruitful collaborations!

Ahmad, Luca and Krzysztof Klimaszewski proof-read parts of my thesis. Ahmad, Łukasz, Mikołaj Chwalisz and Michael Döring listened to my rehearsals. Big thanks go to all of you for your time and valuable feedback!

I thank also Henning Francke, Irene Hercka as well as Alicja and Adrian Fiedler for their help in the final stage of thesis writing.

Last but not least, I want to thank my family. The process of upbringing and education is extremely challenging. I owe you a lot throughout the whole way towards my PhD, starting from my birth, through kindergarten, primary school, secondary school, graduate studies, and ending on my work at TKN. Thank you soooo, soooo much!

Last but not least, I mention my beloved wife and children, who constantly supported me. I particularly felt it during the last months of the intensive work. I will always remember Ula's words: "Ciiii, tatuś pisze doktorat!"¹.

¹"Shush, daddy's writing a PhD thesis!"

Abstract

Today's core networks are permanently powered on and consume non-negligible amount of energy. Traffic varies over time giving an opportunity to switch off or put into standby mode a subset of network devices in order to save energy in low-demand hours. Line cards are targeted to be dynamically switched on and off since their activation and deactivation times are expected to be sufficiently small, their power consumption is significant (400–500 W per line card), and a sufficiently large set of line cards is usually installed in the network. The focus of this thesis is on Energy-Aware Adaptive Routing Solutions (EA-ARSs) which utilize traffic variation, traffic rerouting and sleep modes in order to save energy consumed by line cards in Internet Protocol (IP)-over-Wavelength Division Multiplexing (WDM) networks.

We collect from publicly available sources (product data sheets, research articles, and databases) extensive sets of realistic input data for the EA-ARSs that are crucial for energy saving, namely traffic data and power consumption data of single network devices. Extensive sets of traffic matrices containing traffic demands between all node pairs in the network cover periods of days, weeks, and months with granularity of 5 min, 15 min, a day, and a month. Physical supply topologies corresponding to the traffic matrices are reported too. The power consumption of single network devices determines the amount of power that can be saved in the whole network by switching off subsets of devices.

Using the realistic input data we estimate the potential power and energy saving assuming different levels of freedom of rerouting in an IP-over-WDM network. The proposed approaches Fixed Upper Fixed Lower (FUFL), Dynamic Upper Fixed Lower (DUFL) and Dynamic Upper Dynamic Lower (DUDL) differ with the possibility of rerouting of traffic demands over the logical (IP) topology, and the possibility of rerouting of lightpaths over the physical topology. A Static Base Network (SBN) is used as a starting point for EA-ARSs. The SBN determines devices installed in the network. Sophisticated Mixed-Integer Linear Programming (MILP) formulations are used for the SBN design as well as for DUFL and DUDL computations, which are highly complex optimization problems. FUFL is a simple, fully distributed heuristic. Our results show that flexibility of the IP routing (Dynamic Upper) contributes the most to the energy saving. Additional flexibility of routing in the WDM layer (Dynamic Lower) brings marginal savings. The simple approach FUFL brings significant savings, which are not as high as the savings brought by DUFL though. Furthermore, they depend on the ratio of the maximum total traffic demand and the capacity of a WDM channel. Spatial distribution of traffic demands also influences the energy savings achieved with FUFL.

A set of evaluation criteria for EA-ARSs is proposed in order to determine the approaches which have a chance to be implemented in an operational network of the future. In this context, an adaptive heuristic called Energy Watermark Algorithm (EWA) is proposed. EWA uses network configuration from previous time period in order to calculate new network configuration. This speeds up calculation of network configuration and reduces reconfiguration costs. Furthermore, EWA has a set of parameters, which allows a network operator to find the preferred trade-off between energy

saving and the Quality of Service (QoS) guarantees. EWA is compared with the Least Flow Algorithm (LFA) and the Genetic Algorithm (GA) taking power and energy consumption, reconfigured traffic, and overload ratio as evaluation metrics. The evaluation is performed on unique network scenarios. Differently to related research, past traffic data is used to design the SBN (determining devices to be installed in the network and their configuration), and different future traffic data is used to evaluate EWA, LFA and GA starting from the SBN designed with a modified version of GA. EWA outperformed LFA and achieved results similar to the GA, having lower computation times and smoother changes of power consumption. The overload is marginal, and observed mainly in the initial phase of the evaluation, when EWA could not use the network configuration from previous time period, but started from the overprovisioned configuration with all devices powered on.

Implementation issues including control and management planes are discussed in this thesis as an outlook toward future work. FUFL and DUFL have been demonstrated on an IP-over-Gigabit Ethernet (GbE) testbed using off-the-shelf equipment. No traffic loss is observed.

An extensive survey of related work is provided in this thesis, where a class of EA-ARSs is distinguished from the class of Energy-Aware Network Design (EA-ND) approaches. For both EA-ARSs and EA-ND classes, we distinguish approaches targeting devices in the IP, WDM and in both layers for energy saving.

Zusammenfassung

Heutige Kernnetze sind permanent in Betrieb und verbrauchen erhebliche Mengen an Energie. Der Datenverkehr schwankt über die Zeit. Um in Zeiten geringen Datenverkehrs Energie zu sparen, können Netzwerkkomponente in einen Schlafmodus mit geringem Energieverbrauch oder komplett abgeschaltet werden. Bei näherer Betrachtung der Netzwerkkomponenten verbrauchen Leitungskarten (engl. line cards), die in hoher Anzahl in einem Netzwerk installiert sind, einen signifikanten Anteil an Energie (Leistung von 400–500 W pro line card). Da geringe Aktivierungs- und Deaktivierungszeiten für die Leitungskarten erwartet werden, sollen sie dynamisch ein- und ausgeschaltet werden. Der Schwerpunkt dieser Dissertation liegt in Energy-Aware Adaptive Routing Solutions (EA-ARs), die Verkehrsschwankung, Wegewahl und Schlafmodi anwenden, um durch die Leitungskarten verbrauchte Energie in Internet Protocol (IP)-over-Wavelength Division Multiplexing (WDM) Netzwerken zu sparen.

Wir sammeln aus öffentlich zugänglichen Quellen (Datenblättern der Hersteller, wissenschaftliche Veröffentlichungen, Datenbanken) umfangreiche und für das Energiesparen kritische Eingangsdaten, nämlich Daten über Verkehrsbedarf und Leistungsverbrauch einzelner Netzwerkkomponenten. Die Verkehrsdaten bestimmen die Stunden des niedrigen Datenverkehrs. Umfangreiche Sätze von Verkehrsmatrizen protokollieren den Datenverkehr zwischen allen Knotenpaaren in Zeiträumen von Tagen, Wochen und Monaten mit der Granularität von 5 Minuten, 15 Minuten, einem Tag und einem Monat. Die Netzwerktopologien, die den Verkehrsmatrizen zugrunde liegen, sind ebenfalls in dieser Dissertation dokumentiert. Der Leistungsverbrauch einzelner Netzwerkkomponenten bestimmt den Leistungsverbrauch, der durch Abschalten von Teilmengen der Netzwerkkomponenten gespart werden kann.

Unter Verwendung der realistischen Eingangsdaten wurde das Potenzial der Leistungs- und Energieersparnis abgeschätzt, das sich durch Rerouting mit verschiedenen Freiheitsgraden in einem IP-over-WDM Netzwerk ergibt. Insbesondere unterscheiden sich die vorgeschlagenen Ansätze Fixed Upper Fixed Lower (FUFL), Dynamic Upper Fixed Lower (DUFL) und Dynamic Upper Dynamic Lower (DUDL) in den Möglichkeiten des Reroutings des Verkehrsbedarfs über die logische (IP) Topologie und des Reroutings von Lichtpfaden über die physikalische Topologie. Ein statisches Basisnetzwerk (engl. Static Base Network (SBN)) wird als Startpunkt für EA-ARs benutzt. Das SBN bestimmt die im Netzwerk installierte Netzwerkkomponente. Bei der Bestimmung von SBN als auch bei den DUFL und DUDL Berechnungen wurden komplexe Verfahren zur Optimierung basierend auf einer Mixed-Integer Linear Programming (MILP) eingesetzt. FUFL ist eine einfache, völlig dezentralisierte Heuristik. Unsere Ergebnisse zeigen, dass IP Rerouting (Dynamic Upper) zu der größten Energieersparnis führt. Eine zusätzliche Flexibilität von Routing in der WDM Schicht hat nur eine marginale Energieeinsparung zur Folge. Der einfache Ansatz FUFL bringt wesentliche Energieeinsparung, die jedoch nicht so hoch wie die durch DUFL erbrachte Energieeinsparung ist. Darüber hinaus hängt die Energieersparnis vom Verhältnis zwischen dem maximalen totalen Verkehrsbedarf und der Kapazität eines WDM Kanals ab. Auch die Verteilung des Verkehrsbedarfs beeinflusst die erreichbare Energieersparnis von FUFL.

Ein Satz von Kriterien zur Evaluierung für EA-ARs wurde vorgeschlagen, um die Eignung einer Implementierung von EA-ARs in einem zukünftigen Netzwerk prognostizieren zu können. In diesem Zusammenhang wurde eine adaptive Heuristik namens Energy Watermark Algorithm (EWA) vorgeschlagen. EWA benutzt die Netzwerkkonfiguration aus der vorherigen Zeitperiode, um eine neue Netzwerkkonfiguration zu berechnen. Das beschleunigt die Berechnung der Netzwerkkonfiguration und reduziert Kosten, die bei der Umkonfiguration entstehen. Darüber hinaus verfügt EWA über einen Satz von Parametern, die dem Netzwerkoperator einen Kompromiss zwischen dem Energiesparen und der bevorzugten Dienstgüte ermöglicht. EWA wird mit dem Least Flow Algorithm (LFA) und Genetic Algorithm (GA) verglichen, im Hinblick auf Leistungsverbrauch, Energieverbrauch, Verkehr der umgeleitet werden muss und Überlastquotient. Die Evaluierungsstudie wurde an spezifischen Netzwerkszenarien durchgeführt. Es wurden jeweils verschiedene Verkehrsdaten zur Bestimmung der SBN und zur Evaluierung der betrachteten Algorithmen benutzt. EWA übertraf LFA und hat (mit geringerer Rechenzeit und feineren Änderungen der Leistung) vergleichbare Ergebnisse zum GA erzielt. Die Überlast war marginal und wurde hauptsächlich in der Anfangsphase der Evaluierung beobachtet, als EWA keine Netzwerkkonfiguration aus der vorherigen Zeitperiode benutzen konnte, sondern startete von einer übermäßigen Konfiguration in der alle Geräte eingeschaltet waren.

Implementierungsaspekte inklusive Steuerungs- und Verwaltungsebenen werden in dieser Dissertation im Hinblick auf zukünftige Arbeiten diskutiert. FUFL und DUFL wurden in einer IP-over-Gigabit Ethernet (GbE) Testumgebung mit auf dem Markt erhältlichen Standardgeräten demonstriert. Kein Verkehrsverlust wurde beobachtet.

In der Dissertation wird ein ausführlicher Überblick über die Thematik gegeben. In dem Überblick werden ähnliche Arbeiten besprochen und klassifiziert. Insbesondere werden Ansätze der EA-ARs und des Energy-Aware Network Design (EA-ND) klassifiziert. Für beide Klassen wird analysiert, welche Geräte jeweils in der IP Ebene, der WDM Ebene oder in beiden Ebenen kombiniert ausgeschaltet werden können, um Energie zu sparen.

Table of Contents

Acknowledgements	iii
Abstract	v
Zusammenfassung	vii
Table of Contents	viii
List of Figures	xiii
List of Tables	xvii
List of Acronyms	xix
1 Introduction	1
1.1 Motivation	1
1.2 Scope of the energy saving techniques	2
1.3 Thesis focus and goals	4
1.4 Thesis contributions	5
1.5 Thesis outline	6
2 Background	7
2.1 WDM	7
2.2 IP	8
2.3 IP-over-WDM	10
2.4 Power consumption of single network elements	12
2.5 Dynamics of traffic	13
2.5.1 Overview of the traffic data and network topologies used in related work . .	14
2.5.2 Gathered traffic data originating from measurements and corresponding network topologies	15
3 Related Work	23
3.1 Energy-Aware Network Design	24
3.1.1 Taxonomy	24
3.1.2 IP and WDM layers	25
3.1.3 IP layer	30
3.1.4 WDM layer	33
3.1.5 Concluding remarks and thesis contributions	36

3.2	Energy-Aware Adaptive Routing Solutions	37
3.2.1	Taxonomy	37
3.2.2	WDM layer	38
3.2.3	IP layer	40
3.2.4	IP and WDM layers	46
3.2.5	Concluding remarks and thesis contributions	48
4	Potential energy savings in the IP-over-WDM networks	51
4.1	Design of Static Base Network	51
4.2	Approaches to save energy	54
4.3	Scenarios	58
4.3.1	Topologies and traffic demands	58
4.3.2	Power and CapEx costs	60
4.4	Results	62
4.4.1	Results for the Germany17 network with DFN matrices over one day . . .	63
4.4.2	Varying input parameters	67
4.5	Conclusion	69
5	A step towards bringing energy saving with EA-ARs into reality	71
5.1	Evaluation criteria for EA-ARs	71
5.2	Energy Watermark Algorithm	73
5.2.1	General idea	74
5.2.2	Main steps	74
5.2.3	Differences to related work	75
5.3	Alternative heuristics	76
5.3.1	Least Flow Algorithm	76
5.3.2	Genetic Algorithm	77
5.4	Metrics	78
5.5	Evaluation scenario	80
5.6	Results	81
5.6.1	Time variation	84
5.6.2	Impact of Static Base Network	86
5.6.3	Impact of load variation	87
5.6.4	Sensitivity analysis	88
5.6.5	Power breakdown	89
5.6.6	Computation times and complexity	90
5.6.7	Monetary savings	91
5.7	Conclusion	92
6	Implementability	93
6.1	Definitions required for the management and control with GMPLS	93
6.2	Experimental activities – a survey	95
6.2.1	MiDORi	95
6.2.2	Experiments on the CARISMA testbed	97
6.3	General steps needed for energy saving	98

6.4	Available methods	99
6.4.1	Traffic monitoring and communication of the monitored information to the central PCE	99
6.4.2	Validation of triggering events	100
6.4.3	Calculation of network configuration	100
6.4.4	Communication of new network configuration and the network reconfiguration	100
6.5	Implementation on a testbed	102
6.5.1	Testbed setting	102
6.5.2	Methods	104
6.5.3	Results	106
6.6	Conclusion	108
7	Conclusion	109
7.1	Summary of contributions	109
7.2	Future work	110
A	List of Symbols	113
B	Model variants for DUFL and DUDL	117
C	Details of Energy Watermark Algorithm	119
C.1	Ensuring Routability of Demands	119
C.2	Establishing Necessary Lightpaths	120
C.3	Releasing Unnecessary Lightpaths	122
D	Power consumption of single network elements	125
D.1	IP devices	125
D.2	WDM devices	129
D.3	Other aspects regarding power consumption	133
E	Networks – nodes and physical supply links	135
F	Publications	139
G	Bibliography	141

List of Figures

1.1	General view on the hierarchy of telecommunications networks with exemplary technologies to realize them [1].	2
1.2	Fluctuation of traffic (total demand) over time.	3
2.1	Model of the considered IP-over-WDM network in an exemplary configuration. Colors of lightpaths determine wavelengths used by the lightpaths.	8
2.2	Router model with exemplary line cards and port cards [2].	9
2.3	An example showing different layers (stacked one on the top of the other) in the considered network model apart from the control and management planes. Traffic units are normalized to the capacity of a single lightpath.	11
2.4	State model of a network device (based on [3]).	13
2.5	Physical supply network topologies [4].	20
2.6	Total demand over the whole periods.	21
2.7	Germany17 network and traffic [4].	22
3.1	Publications cited in this chapter – breakdown over time.	23
3.2	Classification of energy saving approaches.	24
4.1	A simple example showing how the approaches FUFL, DUFL and DUDL decrease the number of active line cards. There are five peak demands ($AB = 1.3$, $AC = 2.1$, $AD = 1.2$, $BC = 0.7$, $CD = 0.8$), which decrease in the low-demand hour ($AB = 0.6$, $AC = 0.4$, $AD = 0.9$, $BC = 0.6$, $CD = 0.0$). The physical topology is fixed (solid lines in subfigures (i)–(l), each line corresponds to a single fiber of capacity $B = 3$ wavelengths). The granularity of the logical link capacity is $W = 1$ (normalized capacity, two line cards). In the low-demand hour 2, 6, and 8 line cards are saved with FUFL, DUFL, and DUDL, respectively.	55
4.2	Physical supply network topology, and source traffic distribution. The area of a node represents its emanating demand. The DFN demands (a) are Frankfurt-centralized in contrast to the DWG demands (b). Notice that the New York node is omitted in the Géant figure (d).	58
4.3	Total demand over time in Tbps for the scenario with 3 Tbps maximum total demand, every day of February 2005, and every month of 2004. The DWG matrices show the expected behavior over time with peaks during the week, low traffic on weekends (a), and a slightly rising demand over the year (b). In contrast, the DFN measurements exhibit peaks caused by single demands.	59

4.4	The figures show the power consumption in kilowatt with the three strategies on the Germany17 network, DFN traffic, 1/3/5 Tbps, every 15 minutes on February 15, 2005. The difference between FUFL and DUFL is much larger than the additional benefit of DUDL. Note that the y-axes are not identical.	64
4.5	Logical topologies of the SBN and of the computed solutions with FUFL, DUFL, and DUDL at 05:30 am for the Germany17 network with DFN measurements, maximum total demand 5 Tbps. The color of a link corresponds to the load (high (red) and low (green)). The width of a link refers to its capacity. The size of a node represents its emanating demand.	65
4.6	Average logical link utilization over all active logical links, Germany17, DFN, every 15 minutes. DUDL and DUFL achieve high lightpath utilization in contrast to FUFL.	67
4.7	The results are independent of the network topology. The figures show the power consumption over time with the three approaches on the Géant and Abilene network with measured traffic demands, scaled to 1 Tbps and 5 Tbps of maximum total demand, every 15 minutes of a day.	68
4.8	The results are independent of the time scale. The figures show the power consumption over time with FUFL, DUFL, and DUDL in the Germany17 network, DFN traffic, 3 Tbps, for a month and a year. The meaning of the curves is the same as in the 15-minute Fig. 4.4(b).	69
4.9	Power consumption over time with FUFL, DUFL, and DUDL in the Germany17 network, 3 Tbps, DWG traffic, 15 minutes over one day. The change of the spatial distribution of traffic affects the success of FUFL. Compare with Fig. 4.4(b).	69
5.1	Total demand over time for the traffic data used for EWA evaluation of energy savings (total load per node 300 Gpn).	81
5.2	Power consumption of active line cards and total reconfigured traffic in the Abilene network on 2004/08/27.	85
5.3	Power consumption of active line cards and total reconfigured traffic in the Géant network on 2005/06/10.	85
5.4	SBN Variation: Energy consumption and reconfiguration ratio in the Abilene and Géant networks on 2004/08/27 and 2005/06/10, respectively.	86
5.5	Load Variation: Energy consumption and reconfiguration ratio in the Abilene network on 2004/08/27.	87
5.6	Load Variation: Energy consumption and reconfiguration ratio in the Géant network on 2005/06/10.	87
5.7	EWA Sensitivity Analysis for the Abilene network on 2004/08/27.	88
5.8	EWA Sensitivity Analysis for the Géant network on 2005/06/10.	88
5.9	Power Breakdown for the Abilene and Géant networks on 2004/08/27 and 2005/06/10, respectively.	89
6.1	Steps needed for energy saving (traffic rates in Gbps, IP link capacities in number of lightpaths denoted as λ in the figure).	99
6.2	Base logical topology.	103

6.3	Differentiation between T_M (time period over which the constantly monitored traffic is averaged) and T_L (the time period after which triggering events are validated).	104
6.4	Utilization of x interfaces on a logical link operated by FUFL.	105
6.5	FUFL and DUFL results on the testbed (mind different y-scales in (c) and (d)). . .	107

List of Tables

2.1	(Random) traffic used in the literature.	14
2.2	Network topologies used in the literature. Uncertain data is marked with “?”.	16
2.3	Summary of network scenarios [4].	18
3.1	Energy-Aware Network Design (EA-ND) methods (IP and WDM layers).	29
3.2	Energy-Aware Network Design (EA-ND) methods (IP layer).	32
3.3	Energy-Aware Network Design (EA-ND) methods (WDM layer).	35
3.4	Energy-Aware Adaptive Routing Solutions (EA-ARSs) targeting power saving in WDM layer.	39
3.5	Energy-Aware Adaptive Routing Solutions (EA-ARSs) targeting power saving in IP layer.	43
3.6	EA-ARSs targeting power saving in IP and WDM layers.	47
4.1	Network topologies and traffic [4].	60
4.2	Set of available routers [2, 5].	61
4.3	Cost of optical components and of line cards [2, 5, 6].	62
4.4	Size of the MILPs to compute the SBN in terms of variables and constraints. Simple upper or lower bound constraints for variables are not counted.	63
5.1	Considered networks and traffic data [4].	80
5.2	Line Card Daily Energy Consumption E^{LC} [kWh].	83
5.3	Total Daily Energy Consumption E^{TOT} [kWh].	83
5.4	Reconfiguration Ratio ξ	83
5.5	Overload Ratio ϕ	84
5.6	Computation Times and Complexity.	90
5.7	Yearly Monetary Cost [thousands of Euros].	92
6.1	Current (“A”) and proposed (“S”) bits in the signaling protocol (ADMIN STATUS object of the Resource reSerVation Protocol-Traffic Engineering (RSVP-TE) path/resp messages) [3].	101
6.2	Traffic demands (bidirectional) in the network.	103
A.1	Notation.	113
D.1	Power consumption of line cards in the IP layer.	125
D.2	Power consumption of shelves in the IP layer.	127
D.3	Power consumption of chassis in multi-shelf systems of Cisco CRS-1 Carrier Routing System using 16-Slot Single-Shelf Systems [7].	128

D.4	Power consumption of chassis in multi-shelf systems of Juniper T1600 core routers using Juniper TX Matrix [8].	129
D.5	Power consumption of transponders / muxponders in WDM layer.	129
D.6	Power consumption of regenerators in WDM layer.	130
D.7	Power consumption of WDM terminals in WDM layer.	131
D.8	Power consumption of OLAs in WDM layer.	131
D.9	Power consumption of optical switches in WDM layer.	132
E.1	Network nodes with corresponding locations and node ids for the Géant and Germany17 networks.	135
E.2	Network nodes with corresponding locations and node ids for the Abilene network.	136
E.3	Physical supply links in the Abilene network.	136
E.4	Physical supply links in the Géant network.	137
E.5	Physical supply links in the Germany17 network.	138

List of Acronyms

ALBF	Alcatel-Lucent Bell Labs
AP	Access Point
AP&LL	Additional Power and Link Load
ARIMA	Auto Regressive Integrated Moving Average
AWG	Arrayed Waveguide Grating
BFS	Breadth-First-Search
BGP	Border Gateway Protocol
CapEx	Capital Expenditures
CLI	Command Line Interface
CMOS	Complementary Metal-Oxide-Semiconductor
CPU	Central Processing Unit
CR-LDP	Constraint-based Routed Label Distribution Protocol
DAISIES	Distributed and Adaptive Interface Switch off for Internet Energy Saving
DAP	Digital Application Processor
DLF	Distributed Least Flow
DMP	Distributed Most Power
DNA	Distributed Network Architecture
DCF	Dispersion-Compensating Fiber
DFN	Deutsches ForschungsNetz
DGE	Dynamic Gain Equalizer
DiPP	Disjoint Path Protection
DLP	Dedicated Link Protection
DPP	Dedicated Path Protection

DUDL	Dynamic Upper Dynamic Lower
DUFL	Dynamic Upper Fixed Lower
DWDM	Dense Wavelength Division Multiplexing
DWG	Dwivedi-WaGner
DXC	Digital Cross-Connect
EA	Energy-Aware
EA-ARS	Energy-Aware Adaptive Routing Solution
EA-DPP	Energy-Aware Dedicated Path Protection
EA-DPP-Dif	Energy-Aware Dedicated Path Protection with Differentiation of primary and secondary paths
EA-DPP-MixS	Energy-Aware Dedicated Path Protection with Mixing Secondary with primary paths
EA-ND	Energy-Aware Network Design
EA-RWA	Energy-Aware Routing and Wavelength Assignment
EAR	Energy-Aware Routing
ECMP	Equal-Cost Multi-Path
ECR	Energy Consumption Rating
EDFA	Erbium Doped Fiber Amplifier
ELH	Extended Long Haul
EMS	Element Management System
EO	Electrical-Optical
EOE	Electrical-Optical-Electrical
ERO	Explicit Route Object
ESA-STH	Energy and Survivability Aware Single Time-period Heuristic
ESIR	Energy Saving IP Routing
ESACON	Energy Saving based on Algebraic CONnectivity
ESOL	Energy Saving based on Occurrence of Links
ESTOP	Energy Saving based on TOPology control

EWA	Energy Watermark Algorithm
EXC	Electrical Cross-Connect
FCC	Fabric Card Chassis
FCS	Fabric Card Shelf
FE	Fast Ethernet
FF	First Fit
FIFO	First In First Out
FPGA	Field Programmable Gate Array
FUFL	Fixed Upper Fixed Lower
GA	Genetic Algorithm
GAGD	Genetic Algorithm for Green Design
GbE	Gigabit Ethernet
GMPLS	Generalized MultiProtocol Label Switching
GMT	Greenwich Mean Time
GRiDA	GReen Distributed Algorithm
GUI	Graphical User Interface
ICT	Information and Communication Technology
IETF	Internet Engineering Task Force
ILP	Integer Linear Programming
IP	Internet Protocol
IPCC	IP Control Channel
ISIS	Intermediate System to Intermediate System
ISCOM	Istituto Superiore delle Comunicazioni e delle Tecnologie dell'Informazione
ISP	Internet Service Provider
IT	Information Technology
KSP	K-Shortest Paths
L2	Layer 2

L2SC	Layer 2 Switch Capable
LBC	Load Based Cost
LC	Line Card
LCC	Line Card Chassis
LCP	Least-Cost Path
LCS	Line Card Shelf
LCW	Least Cost Wavelength
LE-I	Less Energy Incremental
LFA	Least Flow Algorithm
LH	Long Haul
LMP	Link Management Protocol
LSC	Lambda Switch Capable
LSP	Label Switched Path
LR	Long Reach
LSA	Link State Advertisement
LSDB	Link State DataBase
LSR	Label Switching Router
LT	Logical Topology
LTE	Long Term Evolution
LUB	Link Utilization Based
MEMS	Micro Electro-Mechanical System
MIB	Management Interface Base
MiDORi	<u>M</u> ulti-(layer, path and resources) <u>D</u> ynamically <u>O</u> ptimized <u>R</u> outing
MILP	Mixed-Integer Linear Programming
MLR	Mixed Line Rate
MLTE	MultiLayer Traffic Engineering
MPA	Most Power Algorithm

MPR	Multi Path Routing
MPLS	MultiProtocol Label Switching
MTBF	Mean Time Between Failures
MUP	Most-Used Path
NIC	Network Interface Card
NMS	Network Management System
NREN	National Research and Education Network
OADM	Optical Add-Drop Multiplexer
OE	Electrical-Optical
OEO	Optical-Electrical-Optical
OLA	Optical Line Amplifier
OLMUP	Ordered-Lightpath Most-Used Path
OpEx	Operational Expenditures
OSPF	Open Shortest Path First
OSPF-TE	Open Shortest Path First-Traffic Engineering
OSS	Operations Support System
OTN	Optical Transport Network
OXC	Optical Cross-Connect
P2MP	Point-to-MultiPoint
PA-RWA	Power-Aware Routing and Wavelength Assignment
PC	Personal Computer
PCE	Path Computation Element
PIC	Physical Interfaces Card
PLIM	Physical Layer Interface Module
PON	Passive Optical Network
QoS	Quality of Service
RAM	Random Access Memory

RDB	Red-Demand-Based
RIB	Routing Information Base
REG	REGenerator
REO-hop	Renewable Energy Optimization hop
RSB	Request Size Based
RSVP	Ressource reSerVation Protocol
RSVP-TE	Ressource reSerVation Protocol-Traffic Engineering
RWA	Routing and Wavelength Assignment
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RX	Receiver
SBN	Static Base Network
SDH	Synchronous Digital Hierarchy
SDN	Software Defined Networking
SFP	Small Form-factor Pluggable
ShPR	Shortest Path Routing
SLP	Shared Link Protection
SLR	Single Line Rate
SNMP	Simple Network Management Protocol
SON	Self Organizing Network
SONET	Synchronous Optical NETworking
SPP	Shared Path Protection
SPR	Single Path Routing
SPT	Shortest Path Tree
SR	Short Reach
Star-SH&ReR	Start Single-Hop and Reroute
T-1	Transmission System 1
TATG	Time-Aware Traffic-Grooming

TCAM	Ternary Content-Addressable Memory
TDB	Total-Demand-Based
TDM	Time Division Multiplexing
TE	Traffic Engineering
TE-LSA	Traffic Engineering Link State Advertisement
TED	Traffic Engineering Database
TLB	Table Lookup Bypass
TLV	Type Length Value
TM	Traffic Matrix
TX	Transmitter
TSP	Transponder
ULH	Ultra Long Haul
UPC	Universitat Politècnica de Catalunya
VLAN	Virtual Local Area Network
VLCM	Virtual Line Card Migration
VOD	Video On Demand
WA	Wavelength Assignment
WD	Working Day
WDM	Wavelength Division Multiplexing
WE	WeekEnd day
WLAN	Wireless Local Area Network
WPA-LR	Weighted Power-Aware Lightpath Routing
WSS	Wavelength Selective Switch
xDSL	x Digital Subscriber Line

1 Introduction

1.1 Motivation

It was a late evening in Volos, Greece. “*Where are we going out for dinner?*”. This is a question frequently asked by many people for many, many years... It was Sofie who asked it this time. Everybody looked at their smartphones. “*Is there any WLAN available?*”. Yes, there was. Since Bart had the smartest smartphone, he started searching for the location of a restaurant suggested by Sofie. He used the energy of his smartphone, the energy of an Access Point (AP) in the hotel, the energy of switches, routers, search engines,... The address was however difficult to find on the web, so we eventually found out that it was easier to ask for the information at the reception desk.

Information and Communication Technologies (ICTs) are changing our lives. They change the way we think. They often make our lives easier, but have some drawbacks. Energy consumption is one of the aspects that has to be taken into account when developing newer and newer ICTs. The ICTs undoubtedly allow us to save fuel needed to commute if there were no Internet banking, online shops, etc., however the increasing amount of information that is exchanged via the Internet enforces increase of energy consumed by the ICTs [9]. Therefore many efforts are currently put into reduction of energy consumed by ICTs.

This work focuses on core networks, which provide high capacity links between Internet Protocol (IP) routers. The edge core IP routers aggregate traffic originating from access and metro networks. A general view on the hierarchy of telecommunications networks is provided in Fig. 1.1.

The hierarchy of telecommunications networks can be coarsely compared to the system of vehicle roads. The core IP links can be compared to highways, which carry (vehicle) traffic originating from estate alleys (access), and passing through city streets and country lanes (metro). Various technologies can be used to realize the networks at different hierarchy levels. E.g., Wavelength Division Multiplexing (WDM) or elastic networks can be used in the core, Synchronous Optical NETWORKing (SONET)/Synchronous Digital Hierarchy (SDH) or Metro Ethernet can be used in the metro level, and Passive Optical Network (PON), x Digital Subscriber Line (xDSL) or wireless technologies such as Long Term Evolution (LTE) can be used in the access. While the access networks are the biggest energy consumer in the hierarchy of networks nowadays, this is expected to change [10]. The access technology moves from xDSLs to PONs offering high bit rates directly to the end users. This will result in the increase of data transmitted and received by the end users, but no significant increase in the number of end users themselves (at least in the developed countries). This process will be reflected in the traffic flowing through the core. Power consumption of single core devices ranges from several Watts up to hundreds of kiloWatts [5, 6, 11]. Despite the numerous devices deployed in the access, the major power fraction is thus expected to move to the core, which justifies our closer look at this part of telecommunications networks. Furthermore, core devices are easier to manage because they are usually owned by a single operator and their number is significantly lower than the number of access devices.

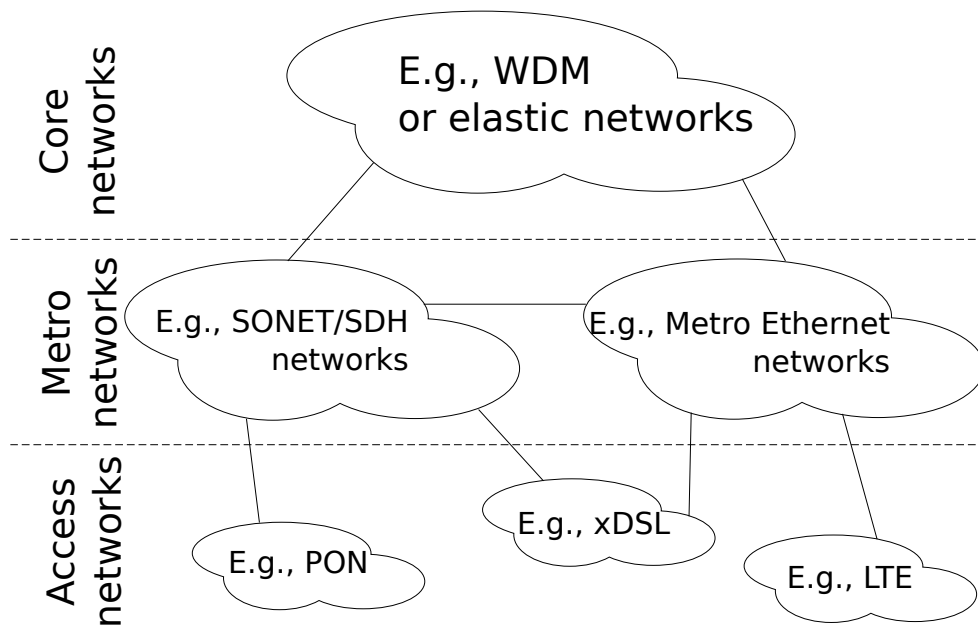


Figure 1.1: General view on the hierarchy of telecommunications networks with exemplary technologies to realize them [1].

1.2 Scope of the energy saving techniques

We draw a general big picture of the various ways of reducing energy consumption of a core network before setting the focus of this work. The various ways of reducing energy consumption of a core network can be classified as follows [12–15]:

1. Power-Efficient Architectures

Although the term “architecture” is used in many different contexts, it is related to the most fundamental decisions in this work. It specifies the number and types of layers in the network, the technologies used in the network (e.g., SDH, WDM, Ethernet), and functions performed by the network. Furthermore, the required level of protection against failures needs to be specified.

Power-efficient architecture concerns also the traffic. Shall different classes of traffic be distinguished? Where can content be placed? Can the content be cached? Policies for routing of the traffic (e.g., is the Multi Path Routing (MPR) allowed?) need to be specified as well.

Translucent IP-over-WDM and transparent IP-over-WDM are two exemplary network architectures. The choice of one or another influences the overall energy consumption of the network [16].

2. Power-Efficient Devices

Energy consumption of a core network is a sum of energy consumed by each device installed in the network. The most straight-forward way to reduce the energy consumption of the network is therefore to improve each device. Vendors of the network devices try to reduce




Figure 1.2: Fluctuation of traffic (total demand) over time.

their power consumption using e.g., improved Central Processing Units (CPUs) with new silicon technologies [12], or new memory technologies like Ternary Content-Addressable Memory (TCAM) [12, 14]. Introduction of sleep modes and making power consumption of a device dependent on load are other ways of reducing the power consumed by a network device. The reduction of power consumption of each device is then reflected in the overall energy consumption of a network.

3. Energy-Aware Network Design

The objective of limited energy consumption of a network can be taken into account at the design stage of the network, where the most power-efficient network devices and routing strategies are considered for installation. The choice of devices that should be installed in the network as well as their configuration should aim at the reduction of energy consumption. The network design strategies may consider usage of sleep modes in low-demand hours. In a similar way, smart geographical placement of network devices can be considered in order to e.g., reduce the energy consumed by cooling devices. A more global perspective would include manufacturing and disposal of network devices, usage of renewable energy, energy spent on maintenance or transmission of energy from the electrical grid to the network devices.

4. Adaptation in high-low demand hours

Traffic varies over time (Fig. 1.2). Core networks transport traffic arriving from metro and access networks. Therefore traffic exchanged between pairs of IP routers in core networks is aggregated, and thus much smoother than traffic in the access. Networks are often dimensioned to carry traffic during peak hours [17]. However, the deployed capacities remain underutilized during low-demand hours, what offers the possibility to save energy by deactivating idle devices. Moreover, the number of idle devices can be increased by dynamic reconfiguration of the network including smart rerouting of traffic demands. Rate adaptation, proxying and virtualization are other ways of adapting the network to the changing traffic in the network [15].

5. Reduction of transmitted data

Putting network devices into sleep mode reduces the power consumption of the network. However, reduced number of network devices usually forces single IP packets to traverse longer paths from their origin to their target. Effectively, the traffic transported by the network increases. Although power consumption of core devices is barely dependent on load nowadays [19, 20], this may change in the future. Therefore, usage of sleep modes should take into account the increase of the traffic traversing the network.

Another way of reducing the amount of traffic in the core network is to reduce it at its origin. The data sent or requested by an end-user sitting in front of his/her computer can be compressed before it is sent out at its origin [12]. Furthermore, each end-user may consider if the traffic that he or she generates is really necessary. This would require however actions beyond the core network or influence of network operator on user data at intermediate stages of transported data.

1.3 Thesis focus and goals

The focus of this work is set on adaptation of transparent IP-over-WDM networks to the changing traffic in order to save energy. The complexity of these networks is related not only to multiple layers (IP and WDM), but also to traffic grooming, routing of traffic over the IP layer, and realization of the topology seen at the IP layer by the WDM layer. Furthermore, the management and control planes are discussed within this thesis, as the mechanisms needed to implement the adaptive energy saving schemes on a real network.

The following research challenges are addressed in this thesis. First, we pose the question about potential energy savings that can be theoretically achieved in the transparent IP-over-WDM networks by switching off line cards in low-demand hours. Second, we estimate savings that can be achieved with heuristic approaches, which are more realistic to be used in real network operation. Eventually, we check the possibility of implementation of the adaptive energy saving schemes looking at the control and management planes.

More in depth, we look at the possibilities of saving energy during low-demand hours, where it is possible to dynamically change the configuration of the network according to the changing traffic demands. There are different types of devices in the network such as Optical Line Amplifiers (OLAs), Optical Cross-Connects (OXC), line cards and IP routers. We focus on the energy consumed by line cards and transponders which are potentially connected to them. The line cards and transponders are responsible for Optical-Electrical-Optical (OEO) and Electrical-Optical-Electrical (EOE) conversion which require high power. At the same time, these devices are expected to have sufficiently small activation/deactivation times for dynamic operation [21].

The idea is to find a balance from the power-efficiency perspective between bypassing the power-hungry line cards and grooming traffic into as few links (lightpaths) as possible, which allows high utilization of the capacities installed in the network. At the same time, a set of constraints needs to be taken into account, just to mention Quality of Service (QoS) and protection against failures.

Network resources are often doubled for (1+1) protection purposes. We do not consider this type of resources in this work, although the proposed algorithms could be used for that part of core networks as well.

We leave aside the choice of network architecture, improvement of the devices themselves, and the reduction of transmitted data, mentioned in Section 1.2. The focus of this thesis is on adaptation of the network in the high-low demand hours, however the Energy-Aware Network Design (EA-ND) is also partially tackled in order to ensure that the energy savings achieved through the sleep modes in low-demand hours are not due to high power-inefficiency of the network designed in power-unaware fashion.

1.4 Thesis contributions

The first main contribution of this thesis is a collection of realistic network scenarios to evaluate energy saving approaches. In particular, we gathered extensive information about power consumption of single network elements, information about realistic network topologies, and traffic originating from measurements, which describe traffic exchanged between all network nodes evolving over time (days, weeks, months). Particularly the information about traffic is critical for the evaluation and extremely difficult to obtain.

The second main contribution of this thesis, is the evaluation of potential of energy saving using Energy-Aware Adaptive Routing Solutions (EA-ARs) considering different levels of rerouting in the IP and WDM layers of an IP-over-WDM network. We found out that a flexibility of rerouting of IP traffic demands over the topology seen at the IP layer (a logical topology) contributes the most to the saving of energy consumed by line cards. Additional flexibility of routing at the WDM layer does not provide significant further savings. No flexibility of rerouting provided in the simple, fully distributed approach called Fixed Upper Fixed Lower (FUFL) still allows for energy saving. The savings achieved with FUFL depend on spatial distribution of traffic demands as well as on the ratio between the total demand and the granularity of capacity of logical links.

The third main contribution of this thesis is a proposal and evaluation of an adaptive heuristic called Energy Watermark Algorithm (EWA). EWA utilizes the freedom of IP routing and the adaptation of logical topology. At the same time, it is adaptive in nature by modifying the current network configuration, and not computing a completely new one. The adaptivity influences EWA computation time, which is important for its live operation. We showed that significant energy savings can be achieved with EWA keeping the reconfiguration costs low. Clear triggering events have been defined so that the network can timely react on the increase or decrease of the traffic finding a balance between energy saving and QoS. The following assumption was made in the EWA evaluation: different (past) traffic data was used to determine network devices installed in the network (network design), and different (future) traffic data was used to determine the devices deactivated in low-demand hours (network operation). This realistic assumption is rarely followed in the related work.

Eventually, two of the proposed approaches FUFL and Dynamic Upper Fixed Lower (DUFL) have been demonstrated on an IP-over-Gigabit Ethernet (GbE) testbed showing that it is possible to remotely manage and control a network in a dynamic manner, so that energy is saved and no traffic is lost. Extensions that need to be made in the management and control planes are discussed in the context of IP-over-WDM networks, focusing mainly on the Generalized MultiProtocol Label Switching (GMPLS) mechanisms.

1.5 Thesis outline

The outline of this work is as follows. Background information including network model is presented in Chapter 2. It provides an overview of the type of network that we look at. It contains the terminology and definitions used throughout the thesis, including the fundamentals of the technologies used in an IP-over-WDM network and interrelations between the layers. Moreover, the key input data for the considered energy saving approaches is discussed, i.e., power and traffic. First, the power consumption data of single network devices is crucial for the estimation of the energy savings that can be achieved with any energy saving approach [5]. Second, sleep modes of network devices can be utilized in low-demand hours. Therefore realistic data describing dynamically changing traffic demands between all node pairs [4] is crucial for any energy saving solution. The traffic data is accompanied with network topology information.

Next, the related work is surveyed in Chapter 3. In particular the Energy-Aware (EA) approaches to network design (EA-ND) and to adaptive usage of sleep modes in low-demand hours (EA-ARs) are presented. Solutions targeting devices in WDM, IP and both layers are distinguished. The aim of this chapter is to provide an overview of the flourishing green backbone networking research, and to identify the contributions of this thesis.

Using the network model and key input data described in Chapter 2, the potential of energy savings is estimated in Chapter 4. Capital Expenditures (CapEx)-minimized network is used as a reference, i.e., as a Static Base Network (SBN). Three different approaches for determining the network elements to be switched off in the low-demand hours are proposed. They differ by the level of freedom of rerouting in IP and WDM layers. Two of the three approaches (Dynamic Upper Dynamic Lower (DUDL) and DUFL) are Mixed-Integer Linear Programming (MILP) problems which require long computation time (in the context of dynamically changing traffic) to find a solution. The third solution FUFL is a simple, quick and fully distributed heuristic.

Differently from the contribution in Chapter 4, the heuristic solution EWA proposed in Chapter 5 explicitly addresses several constraints that an energy saving solution has to face in reality. This includes limited computation time, reconfiguration costs, impact on QoS, and no knowledge of future traffic.

The issues concerning implementability are discussed in Chapter 6. They include not only the time required to compute network configuration, but also the methods for providing input data for the computation and the methods for distributing the computed solution to all network nodes as well as traffic monitoring, events triggering network reconfiguration, and the network reconfiguration itself. Furthermore, related work on experimental activities is surveyed, and two energy saving schemes are demonstrated on an IP-over-GbE testbed.

Chapter 7 concludes the main results of this thesis, and points out open issues for future research. The list of symbols and publications containing parts of this work can be found in the Appendices, together with bibliography and list of publications. Moreover detailed description of mathematical models (DUFL and DUDL) and of an energy saving algorithm EWA as well as source data on power consumption values and considered network topologies can be found in the Appendices.

2 Background

A model of an Internet Protocol (IP)-over-Wavelength Division Multiplexing (WDM) network is presented in order to provide the understanding of the elements that can potentially be switched off to save energy. Their functions and interrelations with each other and with traffic are explained. The WDM layer is presented first. Description of the IP layer follows. Eventually, the multi-layer perspective including traffic is presented. An overview of the network model can be seen in Fig. 2.1 assuming Generalized MultiProtocol Label Switching (GMPLS) as the control plane.

There are two fundamental issues which influence saving energy using adaptation of the network configuration to the changing traffic demands over day and night. The first factor is related to the power consumption of each network element itself, as it determines the amount of power that can be potentially saved by switching this element off or putting it into sleep mode. The second factor is the traffic itself. The shape of the curve showing variation of traffic over time (such as in Fig. 1.2) determines the potential power and energy savings to a major extent. Therefore both power consumption of single network elements and the traffic variability are discussed in this chapter together with the network model.

2.1 WDM

Optical technologies offer high capacities. This is achieved thanks to the Wavelength Division Multiplexing (WDM) technology, where multiple optical signals assigned to a unique wavelength are multiplexed into a single fiber. Each optical signal is generated by an optical transmitter (e.g., semiconductor diode laser or Fabry-Parot laser [22]) and received by a receiver (e.g., PN or PIN photodiodes [22]). Both the transmitter and the receiver are located in a transponder or a colored line card (see Section 2.2).

The optical signals are added or dropped at the network nodes using Optical Cross-Connects (OXC)s.¹ An OXC has the ability to dynamically change its configuration. It can also pass the optical signal from its input to its output. It is composed of arrays of mirrors, which can be adjusted using Micro Electro-Mechanical Systems (MEMSs).

The optical signals that traverse the same fiber installed on a fiber link are multiplexed together. A fiber can carry up to B multiplexed optical signals.² In order to overcome the attenuation, two amplifiers called pre-amplifier and booster are used at the ends of each fiber. The pre-amplifier and booster together with the multiplexer/demultiplexer are often referred to as WDM terminals, which provide the interface between fibers and OXC)s. The optical signal traversing a fiber needs to be amplified every R kms by Optical Line Amplifiers (OLAs). Dynamic Gain Equalizers (DGEs) are placed at some OLAs (e.g., at every fourth OLA [2]) to compensate for any channel power tilt

¹Functionality of Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and OXC)s is very similar. We consistently use the term “OXC)” in this thesis.

²The notation table with symbols used throughout this thesis is available in Appendix A




Figure 2.1: Model of the considered IP-over-WDM network in an exemplary configuration. Colors of lightpaths determine wavelengths used by the lightpaths.

introduced by the OLAs, optical filters, or the Raman effect inside the transmission fiber [2]. We assume that signal regeneration is done in the IP layer, and no dedicated regenerators are used.

The optical WDM channel originating and terminating in the transponders and traversing two or more OXCs is called a lightpath. A lightpath may span multiple fiber links and has capacity W bps. Each intermediate node traversed by the lightpath essentially provides an optical bypass facility [22]. The lightpath can be assigned a unique wavelength on all physical links that it traverses or wavelength converters can be used at intermediate nodes. The assignment of wavelength and choice of the set of physical links that a lightpath traverses is referred to as a Routing and Wavelength Assignment (RWA) problem.

The physical topology consists of physical nodes and fiber links. A fiber link consists of one or more fibers. The physical supply topology (in contrast to the physical topology) determines the nodes and links where network devices and fibers can be installed at the network design stage. The physical topology is the topology, where network devices and fibers have already been installed in the network design stage. All lightpaths between the same pair of nodes can traverse the same or different paths (set of physical links). Each wavelength can be used only once at each fiber.

2.2 IP

The IP is the basic protocol of networking today providing procedures to allow data to traverse multiple interconnected networks [23]. IP routers are installed in the nodes of the network.




Figure 2.2: Router model with exemplary line cards and port cards [2].

A router model proposed in [2] is presented in Fig. 2.2. It consists of a basic node, slot cards and port cards. The basic node includes the control plane and data plane software, switch matrix power supply and cooling, as well as the physical and mechanical assembly. The basic node can be equipped with slot cards and port cards of different types. The port cards are interfaces from/to the lower hierarchy networks and to other nodes. One port card corresponding to each slot card is assumed in this work. A slot card equipped with port card(s) form a line card. The line cards can either be colored or gray [2]. A colored line card generates optical signals, which can be directly fed into the fiber towards the next network node, and therefore does not require any transponder. A gray line card is a Short Reach (SR) interface and requires a transponder, which converts the SR signal into a Long Reach (LR) one. A colored line card or a transponder terminate lightpaths introduced in Section 2.1. We assume the usage of colored line cards in this thesis, since they are more energy-efficient than gray line cards combined with transponders. The advantage of using gray line cards with transponders is the flexibility of choosing solutions of different vendors in the IP and WDM layers.

The core IP routers usually have a modular structure, where the basic node consists of one or more Line Card Shelves (LCSs) interconnected by Fabric Card Shelves (FCSs) [7]. The modular structure allows proper choice of the router according to the required capacity.

Lightpaths are transparent to the IP routers. All parallel lightpaths (regardless of their realization in the WDM layer) between a pair of IP routers constitute a logical link (IP link). The IP routers correspond to logical nodes.

Source-target traffic demands arrive at the IP routers from the lower hierarchy networks (Fig. 1.1). The traffic passes through several steps of aggregation when it goes from the end users through access and metro networks up to the core. The traffic has to be routed through the logical topology in order to get from its source node to its target node. There are different methods of routing of traffic demands. In general, Single Path Routing (SPR) and Multi Path Routing (MPR) can be distinguished. The SPR (such as the one computed with Dijkstra algorithm [24]) is simpler than the MPR (which can be computed using e.g., K-Shortest Paths (KSP) algorithm [24]), since no

decision needs to be taken about what fraction of traffic should take which path. However, the MPR offers the potential of higher power savings, as the traffic can be split over multiple paths so that the logical links are filled to a high extent.

Basic IP routing protocols are Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (ISIS) [24]. Both of them are link state routing protocols. This means that the state of an IP link influences routing of traffic. The state of the link is advertised by the IP router attached to it. The advertisements are broadcasted to all IP routers in the network using flooding or spanning tree, i.e., each router must redistribute all the advertisements so that all routers receive information about all links in the network, and can calculate the same routes over the network.


In a GMPLS controlled network, information about links (called Traffic Engineering (TE) links) is stored in a Traffic Engineering Database (TED). Information about state of the links is exchanged using Link State Advertisements (LSAs) in OSPF. Dynamic changes in traffic routes can be made through modifications in network link metrics [24]. Changes in the link metrics can even result in the increased number of IP hops that traffic will traverse from its source node to its target node, which can however contribute to the lower power consumption of the network. GMPLS mechanisms used for the management and control of the network are described in more details in Chapter 6 (particularly Sections 6.1 and 6.4).

2.3 IP-over-WDM


Looking at both the IP and WDM layers, line cards and OLAs respectively are the devices that can be targeted for switching off in the low-demand hours. The IP routers cannot usually be switched off because of the constantly present traffic originated from or targeted to them. The same way of reasoning makes switching off the whole OXCs difficult. Thanks to the modular structure of the IP routers, some parts of them (i.e., LCSs and FCSs) could be switched off. However, their boot times are expected to be much higher than the times needed to activate and deactivate line cards. Dynamic operation of OLAs is more difficult than dynamic operation of line cards due to transient (thermal) effect of optical transmission [25].

Routing is present in both the IP and the WDM layers. Lightpaths need to be routed through the physical topology in order to provide connectivity between the IP routers. On the top of them, the IP traffic has to be routed over the logical topology in order to get from the source to the target node. Introducing flexibility of routing into one or both layers increases the complexity of the network operation, but offers potential of energy saving.


Fig. 2.3 shows a summarizing example of the network configuration given as input a traffic matrix visualized in the subfigure 2.3(a). Note that the traffic units are normalized to the capacity of a single lightpath. Going down through the layers, the traffic needs to be routed over the logical topology (Fig. 2.3(b)), where each logical link has a certain capacity in terms of number of lightpaths (Fig. 2.3(c)) and corresponding line cards located in IP routers. Traffic demands may be potentially split and traverse multiple logical paths (e.g., routing of the traffic demand between nodes A and D over the logical topology shown in Fig. 2.3(b)). The logical topology needs to be realized over the physical topology. This includes the routing of lightpaths over the physical topology (Fig. 2.3(d)), where each fiber link has a certain capacity in terms of number of fibers (Fig. 2.3(e)). The fibers and their lengths determine the number of OLAs as well as the size of OXCs. Lightpaths connecting the same node pair may potentially traverse different physical paths




(a) Traffic matrix arriving at the IP layer




(b) Routing of traffic demands over the logical topology



(c) Logical topology (number of lightpaths at the IP layer)



(d) Realization of the logical topology over the physical topology (routing of lightpaths)



(e) Physical topology (number of fibers at the WDM layer)

Figure 2.3: An example showing different layers (stacked one on the top of the other) in the considered network model apart from the control and management planes. Traffic units are normalized to the capacity of a single lightpath.

(e.g., lightpaths connecting nodes A and C routed over the physical topology shown in Fig. 2.3(d)).

2.4 Power consumption of single network elements

A network consists of a set of network devices interconnected by fibers at the physical layer. Power consumption of the network can be calculated as a sum of the power consumption of the network devices. Some devices are passive and do not consume power like the fibers themselves. We consider power consumption of single network elements from previous sections in order to identify the most power-hungry ones.

We surveyed the literature for the measurements of power consumption of network elements looking at both research papers and product data sheets. The extensive data is presented in Appendix D based on [5]. We present only a discussion and observations in this section.

There are multiple factors that influence the power consumption to a higher or lower extent. They include configuration of the devices (features and equipment), packet size [19], load [19, 20], and temperature [26].

The following observations have been made based on the power consumption data of the devices in the IP layer:

- Router chassis (i.e., LCSs and FCSs) have the highest power.
- Power consumption of the devices is almost independent of the load [19, 20].
- Power consumption of line cards is significant.
- Power consumption of routers does not linearly scale with the capacity (power consumption versus capacity curves).

The following observations have been made based on the gathered power consumption data of the devices in the WDM layer:

- Power consumption of network elements of the same type varies a lot (e.g., 6–1000 W for OLAs). The reported values are influenced by no consistent measurement procedure (such as Energy Consumption Rating (ECR) [27]) used by vendors. Furthermore, some values assumed in research papers are unreferenced and unjustified (e.g., 1000 W consumed by an OLA in [28]).
- It is difficult to compare the network elements, since their functionalities are not consistently described in data sheets and research papers.

Overall, the devices in the IP layer consume significantly more power than the devices in the WDM layer. Furthermore, the time that is needed to power on a device in the IP layer is significantly lower than in the WDM layer due to the thermal stability [21].

The work [6, 11] led by Ward Van Heddeghem (iMinds) extends [5]. Reference power values were proposed taking the following aspects account: (1) typical values (under typical load and conditions) rather than maximum values are provided; (2) overhead power consumption corresponding to chassis and control equipment is included in the reference values, while the external cooling and facilities (e.g., lighting) is not included; (3) power values correspond to bidirectional




Figure 2.4: State model of a network device (based on [3]).

equipment. The gathered power consumption data is also available in the online database powerlib [29] hosted by iMinds. In this thesis however, we use power values from specific network equipment. These values are not far from the reference values proposed in [6] though. The most important assumption for this work concerns line cards, because they are the network elements which can be activated/deactivated in a short amount of time. We assume 500 W consumed by a line card according to the Cisco CRS-1 router line cards [7]. More precisely, we assume Cisco 4-port 10-GE Tunable WDMPHY Physical Layer Interface Module (PLIM) [30] consuming 150 W and a Modular Services Card consuming 350 W [31]. Assumed power values of all other devices are precisely defined in the corresponding chapters containing results. All of them are based on data listed in Appendix D.

Although the survey in Appendix D shows that sleep mode is not implemented in the devices of today, a third power-state of network devices (additionally to the on and off states, see Fig. 2.4) is desired. Activation of a device from the sleep state is expected to be much faster than booting-up a device [3]. The same applies to the transition from the on state to the sleep state with respect to the transition from the on state to the off state.

Since activation and deactivation of optical devices such as amplifiers are expected to take long time [21], we do not change their states, but keep them always on, what ensures connectivity in the network. Furthermore, we assume that power consumption of a device in a sleep state is negligible.

Eventually, we note for clarity the difference between power and energy in the context of electric circuits. Power determines the rate at which energy is drawn by an electric circuit, and is measured in Watt [W]. The SI unit of energy is Joule [J]. One Watt is equivalent to one Joule per second, and one Joule is equivalent to one Watt times one second. The energy is consistently used to describe operation of the network over a certain period of time (consisting of multiple time points), while the power is used when describing a single point in time. Both energy and power are used as evaluation metrics in this thesis. Only power is used to describe a network device though.

2.5 Dynamics of traffic

Knowing that there are network devices that consume non-negligible amounts of power, we look at the factor that allows to switch them off or put into sleep mode. We indicated in Fig. 1.2 the traffic diurnal pattern, which gives the opportunity to switch off idle devices in low-demand hours. The idea of the Energy-Aware Adaptive Routing Solutions (EA-ARSs) is that the power consumption of the network follows the changing traffic pattern over time. Therefore the traffic data is a critical input for the energy saving studies.

Table 2.1: (Random) traffic used in the literature.

	Type of traffic	Used in
Spatial distribution	Uniform distribution	[32–73]
	non uniform distribution (shown in a table)	[38, 74–78]
	non-uniform traffic matrix (not shown)	[16, 42, 79–81]
	population based model	[82]
	gravity model	[19]
	synthetically generated TM, lognormal distribution [83]	[84–89]
	Unknown distribution	[90, 91]
Time variation	sine-like function	[28, 36, 92–96]
	ARIMA [97]	[67–71, 98, 99]
	random variation	[66]
	Poisson arrival with exponentially distributed holding time	[44, 45, 100, 101]
	Poisson arrival with negative exponential holding time	[59, 102, 103]

2.5.1 Overview of the traffic data and network topologies used in related work

We first look at the input traffic data, because (1) this data is critical for energy saving; (2) it is very difficult to find complete sets of traffic data originating from measurements and describing traffic between all node pairs over a sufficiently long period of time.

The majority of the related work used random or synthetic traffic data as input for their studies as shown in Table 2.1, where we distinguish spatial distribution and temporal variation of traffic. The first one determines traffic demands between node pairs, i.e., the traffic matrix. The second one determines changes of traffic matrices over time.

Looking at the attempts to artificially model the **spatial distribution** of the network (Table 2.1), the gravity and population based models [19, 82] make an attempt to resemble the reality in contrast to the most popular uniformly distributed traffic demands. Looking at the models reflecting **variation of traffic over time**, the sine-like functions are used to model traffic variation in the IP layer, while the Poisson models are used to describe traffic as requests for lightpaths with a certain holding time. Auto Regressive Integrated Moving Average (ARIMA) model [97] is a time series analysis method used to characterize diurnal traffic variation. No details about traffic assumptions are provided in several studies (e.g., [104, 105]).

We found three cases where traffic data originating from measurements (different than the one used in this thesis) was used. First, traffic data originating from measurements is used in [106–109], however the measurements were taken on 1992/01/12 in the Transmission System 1 (T-1) NSFNET network [110]. Second, public and non-public traffic forecasts were mentioned in [111], however it is unclear what the data exactly is due to its non-public nature and the forecasting. Eventually, a similar situation applies to our works [18, 112] and to [113]. More details of scenarios used in

[18] can be found in [114]. The same traffic variation over time was applied to all demands in [18, 112, 113], and also in [28].

The network topologies are much easier to be found concerning the set of nodes and links. It is much more difficult to find data on the capacity of the links and routing of traffic over these links. Table 2.2 contains a summary of network topologies used in the related work as supplementary information to the traffic data. The information whether network links are unidirectional or bidirectional is often unclear. These doubts are marked in the table with a “?”.

The following hierarchical national networks have been used in addition to the ones shown in Table 2.2:

- Germany with 918 nodes in total with 20 nodes and 29 links (bidir.?) in the backbone [111]³
- “a topology similar to the actual one adopted by one of the largest Internet Service Providers (ISPs) in Italy” with 372 nodes and 1436 unidirectional links with 8 nodes and 24 links in the core [28, 61, 62]⁴

Eventually, random network topologies were used in:

- [19] (created with the Brite network generator utilizing the Waxman method).
- [35] (network with 10-35 nodes, unspecified links)
- [38] (network with 16 nodes, unspecified links)
- [65] (network with 100 nodes, unspecified links)
- [60] (hierarchical network with 10 core, 30 edge and 120 aggregation nodes, where core nodes are randomly connected to other core nodes with probability 0.5, each edge node is connected to the two closest core nodes and to another randomly selected edge node, while aggregation nodes are connected to the two closest edge nodes)
- [43] (probability 0.5 of having a link between any two OXCs)
- [59, 92, 93, 102, 103] (nodes uniformly distributed over an area, links added iteratively according to an algorithm)

2.5.2 Gathered traffic data originating from measurements and corresponding network topologies

We present the traffic data measured in three networks (Abilene, Germany¹⁷ and Géant). The data is unique in the sense that it not only originates from measurements, but (1) it also contains complete traffic demands between all node pairs in the network (what is in contrast to measurements of groomed traffic flowing through logical links), and (2) it covers long periods of time, e.g., we have 1 day covered by 288 traffic matrices, each of 5-minute time granularity and 1 month covered by 28 traffic matrices, each of 1-day time granularity; (3) it is publicly available. The data was collected after an extensive search and numerous discussions with various partners. The data sets

³Total number of links is not reported in [111].

⁴373 nodes are reported in [62].

Table 2.2: Network topologies used in the literature. Uncertain data is marked with “?”.

Name	Layer	Nodes	Links	Comment
6N8L-network	WDM	6	8	[33] (bidir.), [34] (bidir.?)
6N7L-network	WDM	6	7	[90] (unidir.?)
8N12L-network	IP	8	12	[91] (unidir)
12N22L-network	IP	12	22	[76] (unidir.?)
AT&T [115]	IP	114	148	[36] (bidir.?), [116] (115 nodes, 296 unidir. links)
Abovenet 6461 [115]	IP	17	74	[19] (dir.)
Abovenet [115]	IP	366	1932	[85, 87, 89] (unidir.), [86, 88] (968 unidir. links)
COST239	WDM	11	26	[32, 44–47] (bidir.), [16, 42, 74] (bidir., single-fiber), [52, 106] (bidir.?), [117, 118] (one node excluded, 10 nodes and 42 unidir. links), [105, 119] (23 bidir.? links)
COST239b	WDM	11	22	[45, 48, 49]
DT [120]	WDM	14	23	[121] bidir., [53] (bidir.?)
Ebone 1755 [115]	IP	18	66	[19] (dir.)
Ebone [115]	IP	159	614	[85, 87, 89] (unidir.), [86, 88] (307 unidir. links), [122] (87 nodes and 161 unidir. links?)
EON	WDM	16	23	[48]
European	WDM	26	49	[77, 78]
European-like	WDM	11	117	[53] (16, bidir.?), [52] (bidir.?), [55] (single fiber per link)
Exodus 3967 [115]	IP	21	74	[19] (dir.)
Exodus [115]	IP	244	1080	[84, 85, 87, 89] (unidir.), [86, 88] (540 unidir. links)
FT [114]	IP	38	72	[112, 113] (bidir.)
Géant2 [120]	WDM	54	52	[121] bidir., [104] with an additional link
Germany50	WDM	50	89	[82] (8 core nodes, bidir.?)
Internet2	IP	9	26	[123–126] (bidir.)
Level3 [115]	IP	63	285	[36, 122] (bidir.?)
NSFNET [22]	WDM	14	21	[16, 32, 44, 67–72, 98, 99, 101, 107, 127, 128] (bidir.), [34, 56–58, 63, 64] (bidir.?, 19 links plotted in [63]), [108, 109] (bidir., 40 unidir. links mentioned in [108], the same topology for the IP layer assumed)
NSFNET14N22L	WDM	14	22	[75] (bidir.?), [42] (bidir., single-fiber)
NSFNET24N43L	WDM	24	43	[38, 56–58] (bidir.?)
Pan-European	WDM	25	46	[59] (bidir.?)
Ring	WDM	4–32	4–32	[43]
Sprint [115]	IP	516	3186	[85] (unidir.), [116] (52 nodes, 296 unidir. links)
Telecom Italia	WDM	31	53	[77]
Telstra [115]	IP	7	18	[19] (dir.)
TID	WDM	30	96	[32] (bidir.)
Tiscali [115]	IP	41	87	[36, 122] (bidir.?)
USNET13N22L	WDM	13	22	[50, 51] (bidir.?)
USNET24N41L	WDM	24	41	[100] (bidir.)
USNET24N43L	WDM	24	43	[45] (bidir.), [34] (bidir.?), [79, 80] (bidir.?)

are publicly available in the SNDlib library [4] hosted by Zuse Institut Berlin. I contributed to [4] by proposing and postprocessing traffic data sets from the Abilene and Géant networks, so that they have the same unified format, the same units, and contain no traffic from a node to itself.

2.5.2.1 Available data sets

The choice of the network scenarios considered in this thesis was determined by the availability of the traffic data covering longer periods of time (at least 1 day) with multiple traffic matrices. The data correspond to the Abilene [129], Géant [130] and Germany17 [131] networks. Even though it was not directly specified in the original traffic data [129–131], we assume that Greenwich Mean Time (GMT) is used.

Abilene The Abilene network (12 nodes and 15 links, Fig. 2.5(a)) is an American network commonly used in the research community [129]. The set of traffic matrices was collected by Yin Zhang and includes 55872 items. Each of the 55872 traffic matrices has granularity of 5 minutes. The traffic matrices cover altogether $55872 \cdot 5 [min] = 279460 [min] = 194 [days]$ from 2004/03/01 to 2004/09/10. However, there are some gaps within this period:

1. 2004/03/15-00:00 – 2004/04/01/-23:55 (5184 missing Traffic Matrices (TMs), 18 days)
2. 2004/04/16-00:00 – 2004/04/21-23:55 (1728 missing TMs, 6 days)
3. 2004/04/29-00:00 – 2004/04/30-23:55 (596 missing TMs, 2 days)
4. 2004/08/20-00:00 – 2004/08/20-23:55 (288 missing TMs, 1 day)

No traffic matrices were available for the periods listed above [129].

Germany17 The German backbone network Germany17 (called Nobel-Germany in [4]) with 17 nodes and 26 links (Fig. 2.7(a)) has been defined as a reference network in the NOBEL project [131].

The traffic data for the Germany17 network was taken from measurements that were performed in the year 2005 in the national research backbone network operated by the German DFN-Verein [132]. The original DFN data consists of the total end-to-end traffic in bytes every 5 minutes over the day 2005/02/15, every day of February 2005, and every month from January 2004 until February 2005 (rather representative data according to partners from DFN). The original traffic matrices were mapped from the original DFN locations to the Germany17 network according to the smallest geographical distances. There is also another version of the German backbone network (with 50 nodes) available in [4], however we leave it out for this study.

Géant The largest network that we investigated (22 nodes and 36 links, Fig. 2.5(b)) is the pan-European research network Géant, which connects European National Research and Education Networks (NRENs) [130]. Other than in the original data, we treated two German nodes as one in Géant network due to location problems.

The set of traffic matrices was collected by Steve Uhlig using Netflow statistics and Border Gateway Protocol (BGP) Routing Information Base (RIB). It includes 11460 traffic matrices. Each of the 11460 traffic matrices has granularity of 15 minutes. The traffic matrices cover altogether $11460 \cdot 15 [min] = 171900 [min] = 119.375 [days]$, from 2005/05/04-15:00 to 2005/08/31-23:45.

Table 2.3: Summary of network scenarios [4].

Name	Nodes	Links	Time granularity	Time horizon	No. of traffic matrices
Abilene	12	15	5 min	6 months	48096
Germany17 (Nobel- Germany [4])	17	26	5 min 1 day 1 month	1 day 1 month 1 year	288 28 14
Géant	22	36	15 min	4 months	11460

Similarly as in the Abilene traffic data, there are some gaps in this period:

1. 2005/05/04-15:00 – 2005/05/04-15:15 (2 empty TMs, 30 min)
2. 2005/05/19-10:00 – 2005/05/19-16:30 (27 empty TMs, 6.75 h)
3. 2005/05/27-17:00 – 2005/05/27-17:30 (3 empty TMs, 45 min)
4. 2005/05/31-15:45 (1 empty TM, 15 min)
5. 2005/05/31-18:45 – 2005/05/31-19:00 (2 empty TMs, 30 min)
6. 2005/06/28-16:45 – 2005/07/03-10:30 (456 empty TMs, 4.75 days)
7. 2005/07/24-05:45 – 2005/07/25-14:30 (128 empty TMs, 32 h)
8. 2005/08/31-08:00 – 2005/08/31-23:45 (64 empty TMs, 16 h)

Differently from the Abilene data, empty traffic matrices are provided for the periods mentioned above in the original data [130]. We drop the first 36 TMs in order to cover whole days, i.e., 119 days from 2005/05/05-00:00 to 2005/08/31-23:45. It is unclear whether no measurements were taken during these periods, or there was no traffic in the network. Furthermore, TMs for the following periods are almost empty (most of the demand values equal to zero):

1. 2005/06/28-16:15 – 2005/06/28-16:30 (2 almost empty TMs, 30 min)
2. 2005/07/03-10:45 (1 almost empty TM, 15 min)
3. 2005/07/24-05:30 (1 almost empty TM, 15 min)

There seems to be a bug in the traffic matrix 2005/05/27-17:45, since the demand sum is 74862010 Mbps while the demand sum for most of the matrices is in the order of 103 to 105 Mbps. The previous three traffic matrices (2005/05/27-17:00 – 2005/05/27-17:30) are empty (see above). Another unusually high demand can be observed between 2005/06/27-11:30 and 2005/06/27-16:15.

Summary of the networks and corresponding traffic data considered in this thesis is presented in Table 2.3 and Appendix E.

2.5.2.2 Brief analysis of the traffic data

We use the total demand defined as sum of all the elements in each traffic matrix to analyze traffic data available in [4]. We leave out the traffic matrices with unusually high demands mentioned in Section 2.5.2.1.

The total demand over time (covered by all considered traffic matrices) is visualized in Figs. 2.6(a) and 2.6(b) for the Abilene and Géant networks, respectively. The first 36 traffic matrices measured on 2005/05/04 in the Géant network are not plotted in Fig. 2.6(b). The peaks on

2005/05/27 and 2005/06/27 are not plotted for visibility, and neglected throughout this thesis. Fig. 2.7(b) shows the total demand for the set of traffic matrices of finest time granularity (5 min) available for the Germany17 network.

The average total demand is equal to 2.84, 43.1 and 4.57 Gbps for the Abilene, Géant and Germany17 networks, respectively. The total demand in the Abilene and Germany17 network is comparable, particularly taking into account the higher number of nodes in the Germany17 network (17 with respect to 12). The total demand in the Géant network is higher even when taking into account that it is the biggest considered network (22 nodes). We report also the median of the total demand, which are 2.58, 41.56 and 4.65 Gbps correspondingly. The absolute values of the traffic demands in all networks are not high with respect to the capacities of a single WDM channel being 10, 40, 100 Gbps and beyond, particularly when taking into account the number of node pairs in the networks (132, 272 and 462 node pairs in the Abilene, Germany17 and Géant networks, respectively⁵). Therefore it is necessary to scale the original traffic matrices up in order to match the capacities of the modern backbone networks. Even the maximum values are relatively small, namely 11.89, 120.13 and 7.19 Gbps for Abilene, Géant and Germany17 networks, respectively.

The traffic variation over day and night influences the possibility to switch off some components in low-demand hours to a high extent. We report the difference between the 97.5-th percentile and the 2.5-th percentile of the total demand instead of the difference between the maximum and minimum in order to filter out the anomalous data and measurement errors. The difference between the 97.5-th percentile and the 2.5-th percentile equals 4.64, 36.88 and 4.78 Gbps for the Abilene, Géant and Germany17 networks, respectively. These absolute values are again relatively small with respect to capacities of today's WDM technologies, however the relative values normalized to the maximum of the total demand show that traffic variation is significant:

- Abilene: $4.64/11.89 \cdot 100\% = 39\%$,
- Géant: $36.88/120.13 \cdot 100\% = 30\%$,
- Germany17: $4.78/7.19 \cdot 100\% = 66\%$.

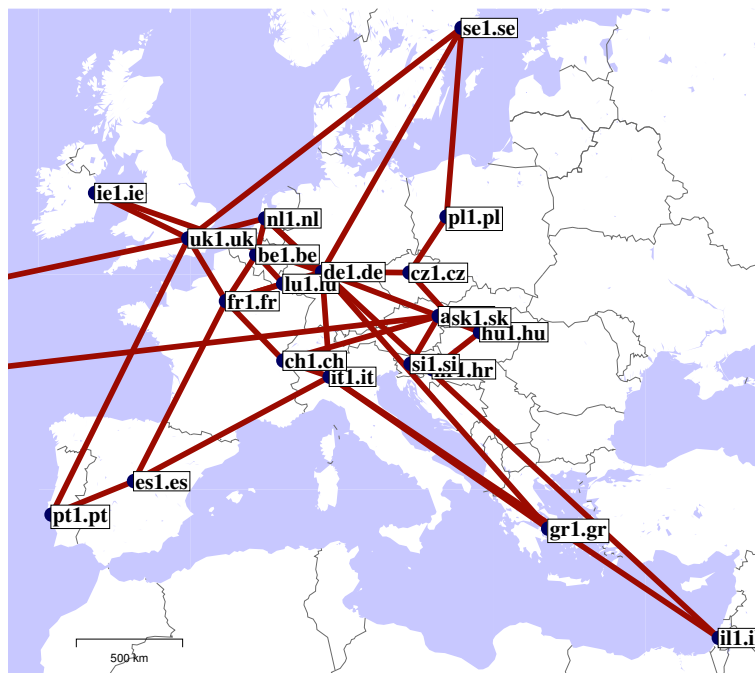
The day-night patterns can be easily recognized in all the curves from Figs. 2.6(a), 2.6(b), and 2.7(b). However, apart from the anomalies reported in Section 2.5.2.1 some spikes can be found in these curves, particularly in the Abilene network. The reason of these spikes is unclear (e.g., measurement errors, peculiar activities within research networks, etc.). A more detailed analysis of single traffic demands instead of the total demand over time (performed within a student project by Olivier Engelbert Eyebe and Mickaël Guth) showed further anomalies. Therefore, representative traffic data without anomalies is chosen to perform optimizations and simulations in this thesis. The data from the SNDlib library [4] was used in [32, 66, 88, 116, 133–135].

Unfortunately, neither access to a commercial backbone network nor up-to-date measurements in such a network were available to us. It is a challenging task for future work. It is challenging for several reasons, i.e., access to the commercial backbone network, amount of data that needs to be captured in order to create a set of traffic matrices with fine granularity over a long period of time, and confidentiality issues. The traffic data presented in this section and available at the SNDlib library [4] is to the best of our knowledge the most complete one and publicly available at the same time.

⁵Traffic demands from a node to itself are neglected [4].




(a) Abilene




(b) Géant (the node ny.ny located in United States is not shown in the figure)

Figure 2.5: Physical supply network topologies [4].

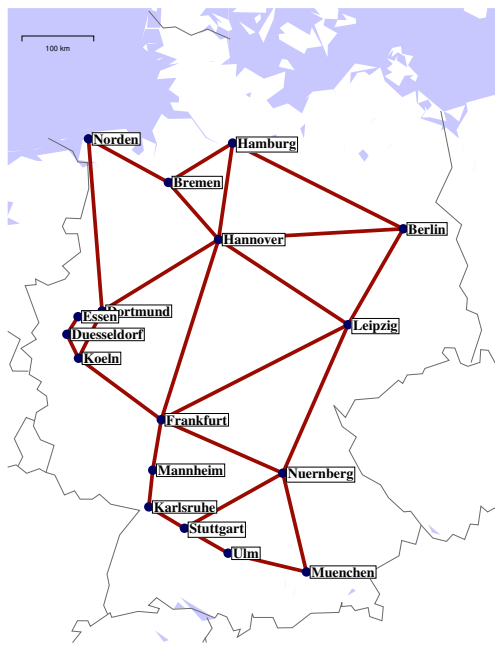


(a) Abilene (194 days from 2004-03-01 at 00:00 until 2004-09-10 at 23:55, granularity 5 min, 55872 directed TMs, scale factor 1)

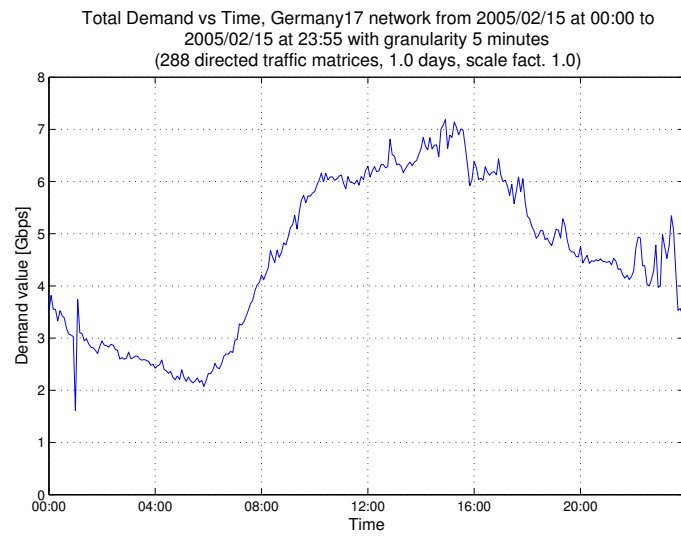


(b) Géant (119 days, 2005/05/05 at 00:00 – 2005/08/31 at 23:45, granularity 15 min, 11424 dir. TMs, scale fact. 1, the peaks on 2005/05/27 and 2005/06/27 not shown for visibility)

Figure 2.6: Total demand over the whole periods.



(a) Physical supply network topology



(b) Total demand over a day with granularity of 5 min

Figure 2.7: Germany17 network and traffic [4].

3 Related Work

Three survey articles on green networking have been published in IEEE Communications Surveys & Tutorials, namely [13–15]. Their submission dates range from August 2009 to March 2010. The last revisions took place between June 2010 and November 2010. Numerous papers on green networking have appeared since then (see Fig. 3.1 for the papers cited in this chapter). Due to




Figure 3.1: Publications cited in this chapter – breakdown over time.

the number of published papers, we limit the scope of this survey with respect to [13–15], and focus on the works related to Internet Protocol (IP)-over-Wavelength Division Multiplexing (WDM) networks in this thesis. The related work is divided into two classes as indicated in Fig. 3.2. The first one tackles the Energy-Aware Network Design (EA-ND), where the devices *to be installed* in the network are determined. The second class is called Energy-Aware Adaptive Routing Solutions (EA-ARSs) which adopt *sleep modes in low-demand hours* of network operation. The term “design” can be used in different contexts. The lack of constraints on installed devices is crucial in this work in order to classify an approach as an EA-ND approach. We distinguish greenfield and non-greenfield approaches. Installation of fibers is considered in the first group in contrast to the second group. The presence of the constraints on installed devices determines EA-ARSs, which focus on network operation in contrast to network design. Moreover, the variation of traffic (over time) must be considered in an EA-ARS. Static routing is allowed though.

The related work is summarized in Tables 3.1–3.6 presented in the following sections. In the cases when there are several publications covering similar energy saving approach, we report all the publications discussing the approach that we are aware of. We use the following notation in the tables in order to keep them compact:

- n/a – not available
- n/c – not considered

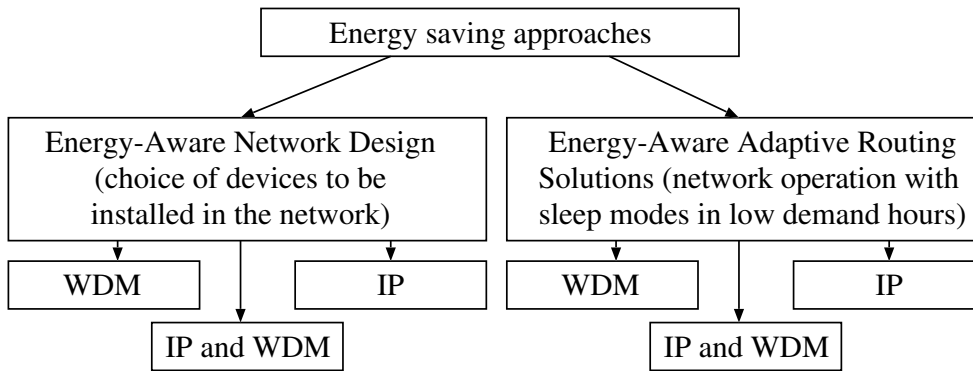


Figure 3.2: Classification of energy saving approaches.

- n/r – not reported
- ppr – postprocessing.

Acronyms are explained in the List of Acronyms (page xix). In order to keep the tables compact, some acronyms (such as Line Card (LC), REGenerator (REG), Transponder (TSP), Logical Topology (LT), Wavelength Assignment (WA), and Traffic Matrix (TM)) are introduced in them, but not used in the text for its readability. Furthermore, the information that was not clearly described in the corresponding paper, but could be guessed from the context was marked with a “?” in the tables available in this chapter. The considered traffic and network scenarios are summarized in Table 2.1 and Table 2.2. Parts of this chapter have been published in [136–138].

3.1 Energy-Aware Network Design

Taxonomy used for the EA-ND methods is presented first. We classify the EA-ND methods according to the layers that they target to save energy. The methods targeting jointly the IP and WDM layers are described first. The methods targeting separately the IP and WDM layers are described next. Please note that an EA-ND method may design both layers, however target energy saving in one layer only (particularly the more power-hungry one).

3.1.1 Taxonomy

The main aim of the network design process is to choose devices **to be installed** in the network. The devices include Optical Cross-Connects (OXC), Optical Line Amplifiers (OLAs), WDM terminals, regenerators and transponders in the WDM layer, and routers (including different chassis/shelves/racks) and line cards in the IP layer. Different terminologies are used in the related work. Short Reach (SR) interfaces, router interfaces, Network Interface Cards (NICs), capacity modules, line modules, Transmitters (TXs), Receivers (RXs), transceivers correspond to line cards and transponders. Switching matrix, switch fabrics and 3D Micro Electro-Mechanical Systems (MEMSs), Arrayed Waveguide Grating (AWG) filters correspond to OXC. Moreover, the related operations performed by the devices are explicitly targeted by some related works in order to decrease the energy consumption (assumption of load-dependent power consumption). The related operations include Optical-Electrical-Optical (OEO) conversion, electrical switching, IP

processing (corresponding to line cards and transponders) as well as optical switching corresponding to OXCs. Furthermore, the greenfield EA-ND approaches determine **installation of fibers**, while the non-greenfield EA-ND approaches assume that fibers are pre-installed in the network.

We distinguish methods of **IP traffic routing** over the logical topology (either Single Path Routing (SPR) with a potential constraint on Shortest Path Routing (ShPR) or Multi Path Routing (MPR)) and **routing of lightpaths** over the physical topology (again either SPR with a potential constraint on ShPR or MPR).

Next, we consider the approach to the wavelength continuity problem, which includes potential usage of **wavelength conversion** and **wavelength assignment**. In general, both of these aspects are complementary and can potentially be done in a postprocessing step in order to decrease the complexity of the network design problem.

Furthermore, we look for constraints ensuring **Quality of Service (QoS)**, which in the network design mainly correspond to the constraint on the maximum level of utilization of network resources (overprovisioning). The overprovisioning can be performed already at the stage of the estimation of the peak traffic that the designed network will handle. Next, we look at the indication of **computation time** or complexity of the given approach. Fast computation is not critical for the network design methods in contrast to the EA-ARSSs. The provided information is purely indicative since it depends on network size, amount of devices that can potentially be installed in the network, implementation and computation platform. Computation time of Mixed-Integer Linear Programming (MILP) problems is labeled as “hours” regardless whether it is reported in the corresponding paper or not. It is possible to stop the MILP solver after a given amount of time, however no limitation of optimization gap can be guaranteed in this way.

Consideration of **protection** is the final aspect of the survey, even though protection is not the focus of this work.

Finally, some of the criteria can be neglected when looking at single layers.

IP layer: The criteria purely related to WDM do not apply. They include installation of fibers, routing of lightpaths in the WDM layer as well as wavelength conversion and assignment.

WDM layer: The criteria purely related to the IP layer do not apply. They include IP routing of traffic demands over the logical topology. However, since the routing of lightpaths over the physical topology is considered, and all parallel lightpaths with the same source and target nodes constitute a logical (IP) link, the design of logical topology is covered indirectly by these network design approaches.

3.1.2 IP and WDM layers

The greenfield approaches are described next. The non-greenfield approaches follow. Results of the survey are summarized in Table 3.1.

3.1.2.1 Summary of the approaches with fiber installation

The Genetic Algorithm for Green Design (GAGD) [138, 139] is an extension of the Genetic Algorithm (GA) proposed in [38] (see Section 3.1.4). GAGD covers the WDM additionally to the IP layer covered by GA. Genetic algorithms use the notion of an individual. The individual is defined as a feasible network configuration in [138, 139]. It is evaluated using power consumption as a fitness function. Individuals are created through random selection of parents, crossover

and mutation. In this way they are expected to find solutions close to optimality, but require less computation time than MILP solvers. Two versions of GAGD (joint and two-step design) are considered. Constraints on SPR and MPR are distinguished. Eventually, the authors of [139] study cases where one or multiple logical topologies are realized on the same physical topology. MILP solutions originating from [140] are reported in [138].

The problem of designing a protected backbone network is formulated as a MILP in [111]. It includes IP routing (both single- and multi-path), logical topology design, realization of lightpaths in the physical layer, and installation of fibers. The problem is solved on a network with 20 nodes and 29 physical links using realistic input data. No heuristic approach is used though.

Shen and Tucker present a network design model in [34]. Their MILP minimizes power consumption of the network, and determines the number of fibers installed on a physical link, established lightpaths and their routing, as well as split routing of IP traffic. Various router configurations are not considered. Optimal solution is achieved only for a small network (6 nodes and 8 physical links) with fibers accommodating 16 wavelengths. Heuristic approaches called “direct bypass” and “multi-hop bypass” are proposed. The power-minimized network is compared with the network designed with the objective of cost minimization (in \$). No single-path IP routing is considered in the MILP. The “multi-hop bypass” heuristic seems to be based on the ShPR (indicated also in [67]).

Design of energy-efficient Mixed Line Rate (MLR) networks is tackled in [74, 75]. The authors formulate three MILP problems for a transparent, translucent and opaque IP-over-WDM networks. The number of fibers on a physical link, established lightpaths (taking physical layer constraints into account) and MPR of IP traffic over the virtual topology are kept as variables. The power consumption of basic nodes (IP routers and OXCs) is not included in the objective function, but the devices to be installed can be determined out of the variables mentioned above. No heuristic approach is studied. Similarly to [34], cost minimization (Capital Expenditures (CapEx) cost normalized to 10G transponder’s cost) is considered as an alternative objective function of the MILP problem.

There are several papers published by the group of prof. Elmirghani from the University of Leeds. We leave out the initial works [101, 141] discussing the concept of anycast routing in optical networks, and focus on the more recent and mature works. Employment of renewable energy sources is considered in [67–70]. The authors propose an Integer Linear Programming (ILP) problem and a heuristic called Renewable Energy Optimization hop (REO-hop) targeting minimization of non-renewable energy consumption. Number of multiplexers/demultiplexers, number of fibers on each physical link, number of wavelength channels on each physical and virtual link, number of ports at each node and the multi-path IP routing are kept as variables in the ILP problem. The REO-hop heuristic is based on the multi-hop bypass heuristic from [34]. Traffic demands are attempted to be routed over the virtual topology so that they traverse the maximum number of nodes using the renewable energy. Shall the attempt be unsuccessful, ShPR for the traffic demand is verified. This is in contrast to the “multi-hop bypass” heuristic, where only the ShPR is verified. The NSFNET network loaded with random traffic demands is used for the evaluation. The authors of [67] extended their work with [71], in which they focused on the physical topology design of IP-over-WDM networks (no constraints of the physical supply topology). Variation of the constraints on nodal degree limit and number of links is the main difference from [67]. Furthermore, Dong et. al. considered network design for symmetric and asymmetric traffic (hot node scenario), and CapEx minimization. Data centers, replication schemes, and transporting renewable energy

to data centers (vs. transporting bits to where renewable energy is) are additionally considered in [99]. Embodied energy, energy-efficient data compression and BitTorrent content distribution in optical networks are additionally tackled in [98]. Caching of content is considered in [72]. Content replication, storage as a service and placement of virtual machines are investigated in [128].

A 2-step MILP is proposed in [142]. The network at the IP layer is designed first. It is then used as input for the design of the network at the optical layer. The variables include: number of wavelength channels between node pairs on the precomputed K-shortest physical paths, multi-path IP routing, number of fibers on physical links, number of add/drop Dense Wavelength Division Multiplexing (DWDM) ports (connecting the Reconfigurable Optical Add-Drop Multiplexer (ROADM) to the IP router and to other nodes), and number of line cards determining router configuration (S2 fabric cards, Line Card Chassis and Fabric Card Chassis) at each node. The proposed 2-step MILP is evaluated on a small-size network (6 nodes and 8 bidirectional links as in [34]) and the NSFNET network (14 nodes and 21 bidirectional links). No heuristic approach is proposed, no differentiation of IP routing schemes is considered, and different optimization objectives are not studied.

The model presented in [104] assumes precomputed routing (corresponding to link-by-link grooming, and end-to-end grooming), and counts the number of necessary devices (routers with line cards, regenerators and fibers with optical amplifiers) according to predefined rules and formulas. No MILP problem is formulated.

3.1.2.2 Summary of the approaches without fiber installation

Circuit- vs. packet-switched network architectures are considered in [32]. All intermediate nodes of a lightpath are bypassed in the circuit-switched architecture, and dedicated and only dedicated lightpaths are established for each traffic demand. The opposite holds for the packet-switched architecture, where no bypass is allowed. The circuit- and packet-switched network architectures are similar to end-to-end grooming and link-by-link grooming from [104], respectively. Five different network topologies are considered in [32] (COST239, NSFNET, Germany17, Géant and TID). CapEx and Operational Expenditures (OpEx) costs are analyzed in the function of traffic load.

The MILP proposed in [106] finds the IP routing, installed line cards (determining chassis and interconnecting fabrics), logical topology and its realization in the WDM layer taking into account physical layer constraints, installed transponders and corresponding slave shelves and racks. The COST239 network topology (11 nodes and 26 links) is used in this study with the assumption of a single-fiber per link. Heuristic approaches, variation of IP routing schemes and alternative optimization objectives are not considered.

The authors of [90] present a model for multi-layer network design. They consider two types of line cards (1 Port OC-48/POS and 4 Port OC-12/POS) with gray interfaces, diversified light-path capacities, multiple chassis configurations, and Routing and Wavelength Assignment (RWA). Due to the complexity of the formulated MILP, the problem is optimally solved on a small network (6 nodes and 7 physical links) with one fiber per physical link and 3 wavelengths per fiber. Traffic demands are randomly generated. Similarly to [106], heuristic approaches and variation of IP routing schemes are not considered. Cost-efficient¹ design is investigated as an alternative optimization objective for the network design.

¹Cost seem to refer to the CapEx cost in [90], but it is not explicitly stated in [90].

Yetginer and Rouskas consider a two-layer architecture, where each node is equipped with an OXC and a Digital Cross-Connect (DXC) [33]. Routing of lightpaths over the physical layer and routing of traffic over the lightpaths are taken as variables in the proposed ILP problem. Following the metric proposed in [19], the authors of [33] define the power consumption of the network as a weighted sum of the number of lightpaths and total amount of electronically routed traffic. Three objective functions originate out of this metric: Minimum Number of Active Router Ports (minL), Minimum Amount of Electronically Switched Traffic (minT) and Minimum Power Consumption (minP). A 6-node network with 8 links under random traffic (uniform distribution with varied average value) is studied. The power consumption of a lightpath under no load (fixed power consumption) is equal to 0.25 of the power consumption of a lightpath under full load (maximum power consumption). The results indicate that minP uses only a few more lightpaths than minL. The difference between minP and minT in terms of electronically switched traffic vanishes as the network load is increased.

Two approaches are taken in [100]. First, static traffic grooming problem (given traffic in a predefined set of time slots) is formulated as a MILP problem. Second, dynamic traffic grooming problem (without traffic knowledge, but considering traffic demands arriving one at a time) is tackled using the proposed Time-Aware Traffic-Grooming (TATG) algorithm. The TATG algorithm uses a Grooming Graph with lightpath edges, wavelength edges, and transceiver edges, which are assigned specific weights. The shortest route over the Grooming Graph determines whether a new lightpath should be established for the arriving traffic demand and/or which existing lightpaths shall the demand be groomed into. USNET24N41L network (24 nodes and 41 physical links) is used for evaluation of the MILP and TATG Traffic demands are defined as OC-3, OC-12, OC-48 and OC-192 requests. The wavelength assignment problem is solved by both TATG and the MILP. Even though time-awareness is targeted by the proposed approaches, no constraint on the installed devices is considered (there are always transceiver edges available).

Network design considering different transport architectures was studied in [16, 143]. The architectures (basic IP-over-WDM, IP over Synchronous Digital Hierarchy (SDH) over WDM, transparent IP-over-WDM and translucent IP-over-WDM) differ with the possibility of grooming, regeneration, and additional SDH layer between IP and WDM. Network design problem for all transport architectures was formulated using MILP. Evaluation was performed on the COST239 network (11 nodes and 26 bidir. physical links) and the NSFNET network (14 nodes and 21 bidir. physical links) assuming a single fiber on each physical link.

In [144] the authors focus on two network architectures (IP-over-WDM with gray interfaces and IP-over-Optical Transport Network (OTN)-WDM). The optimization model is not explicitly presented, but the authors explain that optimal router basic node and Electrical Cross-Connect (EXC) basic node as well as the degree of the OXC are selected as a result of the optimization. Underlying fiber topology is given (Germany17 with 17 nodes and 26 links). No heuristic approaches are investigated and the IP routing is not varied. Power-minimized network is compared against the CapEx-minimized one.

Table 3.1: Energy-Aware Network Design (EA-ND) methods (IP and WDM layers).

Method's Name	Devices to install	Installation of fibers	Routing IP	Routing WDM	Wavelength conversion	Wavelength assign.	QoS constraint	Computation time	Protection considered
GAGD (joint and two-step design) [138, 139], Bianco	routers, LCs (lightpaths), OLAs & WDM terminals (fibers)	yes	SPR, MPR	MPR	n/c	n/c	over-prov.	mins.	no
MILP [111], Betker	routers, LCs & TSPs (lightpaths), OXCs, OLAs & WDM terminals (fibers)	yes	SPR, MPR	MPR	n/c	n/c	over-prov.	hours	link disjoint paths
MILP, direct bypass, multi-hop bypass [34], Shen & Tucker	LCs & TSPs (lightpaths), OLAs	yes	SPR (direct bypass, multi-hop bypass?) and MPR (MILP)	MPR	full	n/c	n/c	hours	no
MILPs for opaque, transparent, and translucent MLR networks [74, 75], Chowdhury	routers with electronic processing, LCs & TSPs (lightpaths), REGs, OLAs	yes	MPR	MPR	full (transparent, translucent) or no (opaque)	consid.	n/c	hours	no
ILP and heuristic REO-hop for networks with renewable energy sources [67–70], Dong	routers, ports & TSPs (lightpaths), OLAs, MUXs/DEMUXs	yes	MPR (ILP), SPR (REO-hop)	MPR (ILP), SPR (REO-hop)	full	n/c	n/c	hours	no
2-step MILP [142], Wang	routers, LCs & TSPs (lightpaths), ROADMs, DWDM ports	yes	MPR	MPR	n/c	n/c	n/c	n/r	no
mathematical model [104], Van Heddeghem	routers with LCs & REGs, OLAs	yes	SPR (shortest cycle)	SPR (shortest cycle)	n/c	n/c	n/c	n/r	1 + 1
GA and single-layer ILP [32], Bianco	routers, LCs, TSPs (lightpaths), OXCs, OLAs & WDM terminals	no	SPR (GA), MPR? (ILP)	n/c	no	n/c	n/c	hours (ILP), minutes (GA)	no
MILP [106], Rizzelli	routers, LCs & TSPs, REGs	no	MPR	MPR	no	consid.	n/c	hours	no
MILP [90], Shen	router chassis, LCs & TSPs (lightpaths)	no	MPR	MPR	n/r	consid.	over-prov.	hours	no
MILP [33], Yetginer	router ports with processed traffic, EO and OE conversion	no	MPR	MPR	n/r	consid.	over-prov.	hours	no
MILP and simulation with TATG [100], Zhang	LCs (fixed and traffic dependent parts)	no	SPR, Dijkstra for TATG	MPR	no	consid.	n/c^2	hours (MILP), n/r (TATG)	no
MILP for basic, transparent and translucent networks [16, 143], Musumeci	OXC, REGs, TSPs, LCs (SR interfaces) with IP processing,	no	SPR	MPR	At 3R REGs	consid.	n/c	hours	no
MILP [144], Palkopolou	routers, port cards, TSPs, OXCs	no	n/r	n/r	full	n/c	n/c	hours	no

²Avg. number of hops and blocking probability are analyzed in [100].

3.1.3 IP layer

We found surprisingly few works addressing the EA-ND with the target of energy-efficient IP layer (see Table 3.2).

Chabarek et al. in [19] formulate a MILP problem, which determines the router chassis and line cards that need to be allocated to each node, as well as the IP routing. The problem was solved for 7 logical topologies (7-21 nodes and 18-134 links) under different loads. The authors support their approach with measurements of power consumption of routers (Cisco GSR 12008 and Cisco 7507). They point out that the power consumed by the routers shows only little dependency on the load (data rate, packet size, packet inter-arrival times). Despite router chassis being the most power-hungry elements of the router, its power consumption is highly dependent on the number and type of installed line cards.

The heuristics proposed in [34] have been extended in [35], where the authors propose the Start Single-Hop and Reroute (Star-SH&ReR) heuristic. The physical layer is left aside with the justification that the WDM devices consume significantly less power than the IP ones. Consequently, a simplification has been made that any logical topology can be realized by the resources in the physical layer, what reduces the complexity of the formulated ILP problem. The basic idea of Star-SH&ReR is to start considering the solution given by the single-hop strategy (i.e., a virtual link for each traffic demand) and then trying to modify link weights and iteratively reroute traffic demands so that unnecessary links can be switched off. Additional Power and Link Load (AP&LL) is used to determine link weights. The output of the network (virtual topology) design consists of IP routing (single path) and the installed line cards, which in turn determine the router chassis.

Logical topology design using a MILP, a greedy heuristic Less Energy Incremental (LE-I), and a meta-heuristic GA were studied in [73] and [38], respectively. LE-I starts from a void logical topology and each traffic demand is satisfied by choosing the less power consuming alternative between adding a new direct link or by using an already available path which has enough bandwidth. Three criteria for ordering traffic demands are considered: increasing, decreasing and random order (according to their rates in bps). The individual in GA is defined as a logical topology [38]. It is evaluated using power consumption as a fitness function. OpEx and CapEx are analyzed in [73] with respect to the ratio between power used by an optical transceiver operating at a given rate and the power consumed by a node to process information in the electronic domain at the same rate.

Puype et al. in [39–41] investigate MultiLayer Traffic Engineering (MLTE) with the objective of reducing the power consumption in IP over optical networks. They assign higher routing costs to lightly loaded logical links in order to empty them after IP rerouting. Empty logical links are removed from the actual logical topology to save energy. Routing of lightpaths constituting logical links is not considered. Applying an algorithmic approach to a 14-node network under a traffic pattern based on a uniform distribution, the authors show significant power savings against a full mesh logical topology. Traffic characteristics cover diurnal traffic variations - off-hour traffic is equal to 0.25 of the peak traffic, however no constraint on the installed devices is used. No constraints on the installed devices are reported though. The authors investigate the case when logical topology updates are slower than the diurnal traffic variations and the opposite one under the assumption that power of router interfaces (line cards) depends on the carried traffic. The authors discuss two ways to influence the power vs. bandwidth (load) curve, i.e., idle power reduction (by e.g., matching line rates with traffic volume, or reducing clock rates) and the scaling of equipment power requirements using newer Complementary Metal-Oxide-Semiconductor (CMOS) technol-

ogy. Both of them influence the results of the considered traffic engineering approach. Power values normalized to the maximum power are reported.

The authors of [76] search for a tradeoff between power and performance (defined as the mean delay). They formulate the network design problem with a simple MILP, and consider three objective functions (power objective, performance objective and trade-off objective). However, there is a mistake in the performance objective (wrong units of mean delay in Eq. (7) of [76]), and consequently also in the trade-off objective. A topology with 12 nodes and 22 links was used for evaluation.

Table 3.2: Energy-Aware Network Design (EA-ND) methods (IP layer).

Method's Name	Devices to install	Installation of fibers	Routing IP	Routing WDM	Wavelength conversion	Wavelength assign.	QoS constraint	Computation time	Protection considered
MILP [19], Chabarek	router chassis, LCs	n/a	MPR	n/a	n/a	n/a	n/c	hours	no
Star-SH&ReR [35], Coiro	LCs	n/a	SPR, ShPR	n/a	n/a	n/a	n/c	n/r	no
MILP, LE-I, and GA [38, 73], Ahmad, Bianco & Bonetto	transmitters and receivers	n/a	SPR (Dijkstra in LE-I and GA) and MPR (MILP)	n/a	n/a	n/a	n/c^3	>24 hours (MILP), n/r (LE-I, GA)	no
MLTE [39–41], Puype	router interfaces	n/a	SPR, ShPR with differently weighted links	n/a	n/a	n/a	n/c	hours	no
MILP [76], Sansò & Mellah	routers	n/a	MPR	n/a	n/a	n/a	consid. ⁴	hours	no ⁵

³Number of hops analyzed in [38].

⁴A MILP objective trading power against performance is considered in [76]. However, there is a mistake in the equation defining the objective (wrong units of mean delay in Eq. (7) of [76]).

⁵Lost traffic upon failure and other metrics are analyzed in [76].

3.1.4 WDM layer

Similarly to Section 3.1.2, the greenfield approaches are described first. The non-greenfield approaches follow. Results of the survey are summarized in Table 3.3.

3.1.4.1 Summary of the approaches with fiber installation

We found only two works that considered installation of fibers (greenfield approaches). First, dimensioning of a pan-European optical network based on the current traffic demands, as well as projections over five and ten years is performed in [118]. The output of the dimensioning process is the number of fibers per link and wavelengths per fiber that need to be installed to serve the input traffic matrix, as well as the dimensions of optical switches required to transparently route the input traffic. The ILP problem minimizes the cost of the dimensioned network (in cost units) and not explicitly the power (even though power can be considered as cost). Power consumption of the designed network is performed in postprocessing summing power consumption of all devices installed in the network.

Second, the number of fibers is an unconstrained variable in [81]. The authors of [81] focus on the node architectures differing in the flexibility of traffic grooming. Apart from the ILP formulation minimizing the number of line cards (line modules), an extensive mathematical model calculating the number of other modules (transponder modules, control modules, etc.) is provided.

3.1.4.2 Summary of the approaches without fiber installation

An auxiliary graph is used in [79, 80] to provision WDM networks. Additionally to the power consumption of network devices (transponders, OLAs and WDM terminals), power related to electronic switching, optical switching and Electrical-Optical (OE) and Electrical-Optical (EO) conversion is considered.

The Power-Aware RWA problem is proposed in [43] and formulated as a MILP problem. Energy savings can be achieved by switching off OXCs and optical amplifiers according to three proposed algorithms called Least-Cost Path (LCP), Most-Used Path (MUP) and Ordered-Lightpath Most-Used Path (OLMUP). All the algorithms start with a void network and iteratively add OXCs and lightpaths. They use SPR to route lightpaths, however different weights are assigned to physical links for LCP, MUP and OLMUP. OLMUP consists of a Lightpath Selection phase and a Routing Update phase. “A rough lower bound” is presented in [43] additionally to the results achieved with the proposed algorithms. Investigations of bidirectional rings and generic meshes of up to 32 nodes without wavelength conversion, and with a large number of fibers on each link and of wavelengths per fiber compared to the number of lightpaths (random requests with no dependency on time) revealed that smart routing of lightpaths in the WDM layer may bring significant energy savings against the SPR and First Fit (FF) wavelength assignment. The authors assume that both the power of an amplifier and the power of an OXC are equal to 1 kW, however no justification for this value is given.

The contribution of [77, 78] is two-fold. First, transmission properties of optical transport networks (such as dispersion) are exploited with a split-step method in order to reduce the number of regenerators. Second, wavelength assignment is performed using OPNET SP Guru Transport planner (black-box) so that some physical links can be powered off. The authors consider a European network (26 nodes and 46 links) and Telecom Italia network (31 nodes and 53 links) with and

without OXCs, and scale a traffic matrix to mimic changing demands. The power consumption of an optical amplifier is estimated to 200-500 W and the power of a 3R regenerator to 2-5 kW (unreferenced). The results show that transmission optimization may lead to the elimination of in-line 3R regenerators, and that traffic grooming allows to switch off significant number of links.

Protection in network design is the focus of [42]. The following four protection strategies are distinguished: Shared Link Protection (SLP), Shared Path Protection (SPP), Dedicated Link Protection (DLP) and Dedicated Path Protection (DPP). Four network design problems corresponding to the four protection strategies are formulated as ILP problems. Apart from the power consumed by transponders and OLAs, also power consumption due to traffic dependent and traffic independent electronic processing is considered. Evaluation is performed on the NSFNET14N22L (the version with 14 nodes and 22 single fiber bidirectional links).

The works [46] and [47] focus on the design of the WDM network taking into account protection requirements. Both of them assume a single-fiber topology with a given number of WDM channels, and are given a set of lightpath requests as input. Dedicated protection resources with sleep modes are considered in [46]. Shared Backup Paths and joint minimization of capacity and energy are considered in [47].

Eventually, the focus of [105, 119] is on providing “virtual infrastructure” for Information Technology (IT) servers. The formulated MILP problem is solved for the COST239 network (11 nodes and 23 links) assuming a single fiber on each physical link. 1:1 protection is considered in [105].

Table 3.3: Energy-Aware Network Design (EA-ND) methods (WDM layer).

Method's Name	Devices to install	Installation of fibers	Routing IP	Routing WDM	Wavelength conversion	Wavelength assign.	QoS constraint	Computation time	Protection considered
ILP [117, 118], Tzanakaki	transceivers, OLAs, optical switches	yes	n/a	MPR	no	ppr, median wavelength	n/c	hours	no
ILP [81], Morais	line modules	yes	n/a	SPR, ShPR	n/c	n/c	n/c	hours	no
Power-Aware Provisioning with Auxiliary Graph [79, 80], Xia	OE and EO conver., TSPs, OLAs, WDM terminals minding el. switching, op. switching	no	n/a	SPR, ShPR over auxiliary graph	Allowed at OEO	n/c	no ⁶	n/r	no
LCP, MUP, OLMUP [43], Wu	OLAs & OXCs	no	n/a	SPR	no	FF	n/c	n/r	n/c
split-step method and OPNET SP Guru Transport planner (black-box) [77, 78], Silvestri	REGs, OLAs	n/r	n/a	n/r	n/r	OPNET	n/c	n/r	1+1 DiPP
ILP [42], Musumeci	OLAs and TSPs with electronic processing	no	n/a	SPR	n/c	n/c	n/c	hours	SLP, SPP, DLP, DPP
ILP [46], Muhammad	transmitter, receivers, switching devices, OLAs (fixed and proportional power)	no	n/a	MPR	full	n/c	n/c	hours	1:1 dedicated
ILP [47], Çavdar	transceivers, OLAs, 3D MEMSs, wavelength converters	no	n/a	MPR	yes	n/c	n/c	hours	shared
MILP [105, 119], Tzanakaki	TSPs, OLAs, OXCs, IT servers	no	n/a	MPR	n/r	n/r	n/c	hours	1:1

⁶Blocking probability analyzed in [79, 80].

3.1.5 Concluding remarks and thesis contributions

ILP and MILP are the most frequently used methods for the EA-ND. Solving the ILP/MILP problems is time consuming, however computation time is not critical at the network design stage. A mathematical model was used only in [104] and [39–41]. There are relatively many works covering both the IP and WDM layers. Most of the EA-ND approaches cover equipment related to lightpaths (line cards). Installation of fibers is mostly not considered, particularly when focusing on the WDM layer only. MPR is a common assumption for both the IP and WDM routing, since it reduces complexity of the formulated optimization problems. Wavelength conversion/assignment is considered by several works. It is usually an integral part of the formulated MILP problems. Network overprovisioning is the most frequently used method to address the QoS constraint. Only the authors of [76] address the mean traffic delay in the optimization objective. There are only few approaches covering protection:

1. **IP and WDM layer:** the MILP [111] by Betker et al. and mathematical model by [104] by Van Heddeghem et al.
2. **IP layer:** no related work
3. **WDM layer:** split-step method and OPNET SP Guru Transport planner (black-box) [77, 78] by Silvestri et al., the ILPs [42] by Musumeci et al., [46] by Muhammad et al., [47] by Çavdar et al., and the MILP [105, 119] by Tzanakaki et al.

The thesis contains a model to design a Static Base Network (SBN) with the objective of minimizing CapEx cost (Section 4.1). This model can be applied also to design a network with minimization of power consumption replacing CapEx cost with power cost. In order to limit the length of this thesis, the following contributions from [138, 140] were skipped:

- MILP formulation of the network design problem installing devices in both the IP and WDM layers (routers, line cards and transponders, OLAs and WDM terminals, fibers) with the constraint of SPR
- solutions of the formulated MILP problems provided by a high performance cluster composed of 128 Central Processing Unit (CPU) cores with 568 GB of total memory [145]
- comparison of the network design with the objectives of CapEx minimization and power minimization – solutions of the MILP problems with the two objectives were similar, since for today’s network devices power barely depends on load. Both power and CapEx cost scale with their capacity. This may change in the future though.
- comparison of the network design under the constraints of SPR and MPR of traffic demands in the IP layer – the constraint of SPR still allows the design of energy-efficient IP-over-WDM networks.

The focus of this thesis is on EA-ARs.

3.2 Energy-Aware Adaptive Routing Solutions

Differently from Section 3.1, we start the review of EA-ARSs from the WDM layer, cover the IP layer next, and eventually summarize the approaches covering both the IP and WDM layers. Moreover, we focus on the algorithmic approaches, because solving MILP problems is too time consuming for dynamic network operation, where traffic varies over time.

3.2.1 Taxonomy

There are fewer decisions (variables) to make during network operation than in the network design. However, there are more constraints with respect to the network design. This is reflected in the taxonomy. The main criterion for classification are the **devices targeted for switching on/off**. Second, the **dynamics of routing** is distinguished (either in the WDM layer or in the IP layer or in both layers). Similarly as in Section 3.1.1, **routing in the WDM layer** and **routing in the IP layer** are reported. The main difference to the network design is the requirement to mind the **constraints on the installed devices**, i.e., fibers, OLAs and other network devices cannot be dynamically added (installed) in the network during its operation. On the other hand, energy saving is usually traded with **QoS**. Its consideration by the EA-ARS is another item in the taxonomy. Furthermore, the **computation time** of the energy-efficient network configuration is crucial in the context of dynamically changing traffic. We focus on network topologies of up to 40 nodes in case results on topologies with more nodes are reported. Larger topologies most probably include metro/access networks additionally to the core network. Timely launching of an EA-ARS is crucial. Therefore we report the **triggering events** that are responsible for this. **Operation** of the network can be either centralized or distributed. Any EA-ARS can be centrally operated assuming that there are mechanisms delivering input information, and control mechanisms distributing the computed solution. Distributed solutions are more demanding. Therefore we report centralized network operation of each EA-ARS unless distributed operation is clearly indicated in the corresponding publication. **Network knowledge** is related to the operation. It determines the information that is required by the EA-ARSs being classified in general as local or global. It is obvious that it is easier and quicker to run an EA-ARS that requires only the knowledge about the traffic load on local links than the one that requires a full traffic matrix. Apart from the QoS, the energy saving is traded also with the level of **protection** against network failures. Therefore protection is also considered in this survey, even though many operators run two networks in parallel [114], and the EA-ARSs may be applied to the secondary network only, while the primary one is operated at full speed for protection purposes. While adapting the network to the changing traffic conditions may bring energy savings, the network needs to be reconfigured between consecutive time periods. Therefore it is important to keep the **reconfiguration cost** (in terms of added/deleted lightpaths or rerouted traffic) to a limited extent. While theoretical studies can show high potentials of energy savings, the **knowledge of future traffic** is crucial in order to achieve the savings. Although it is obvious that nobody knows future traffic, many works neglect the restriction that is related to the (short-term) traffic forecast. Eventually, the **control mechanisms** to deliver the required input information for the EA-ARS as well as to perform the actual reconfiguration are necessary.

Some of the criteria can be neglected when looking at single layers.

IP layer: Routing mechanism and its dynamics in the WDM layer, as well as the physical layer constraints are not considered by the EA-ARS working only in the IP layer.

WDM layer: Routing mechanism and its dynamics in the IP layer are not considered by the EA-ARS working only in the WDM layer.

Furthermore, different targeted devices, QoS metrics, triggering events, reconfigurations costs are applicable in the two layers.

3.2.2 WDM layer

EA-ARSs targeting energy saving in the WDM layer are shown in Table 3.4.

Power-Aware Routing and Wavelength Assignment (PA-RWA) [59, 102, 103], represents an enhancement of the well known strategies for lightpath route computation and wavelength assignment in the classical WDM networks. Differently from [43], it includes constraints on the number of fibers (and wavelengths) installed on each physical link. In the first step of PA-RWA called Load Based Cost (LBC), a shortest path algorithm solves the routing subproblem using a dynamic load-dependent function to associate a weight to each optical fiber link. In the second step Least Cost Wavelength (LCW), wavelength assignment is performed considering the load state and the power consumption of the OLAs deployed along the fibers. LCW assigns a cost to each wavelength available on the links belonging to the path computed by LBC, and chooses the wavelength minimizing the cost on the whole path.

A method iteratively switching off fibers in the network is considered in [92, 93]. Various criteria of ordering fibers that are attempted to be deactivated are considered. All lightpaths traversing the fiber to be switched off are attempted to be rerouted. If both working and protection paths for a lightpath are found, all OLAs on the fiber are switched off.

RWA-Bill problem is formulated as an ILP in [50, 51]. The focus of this work is on the energy bills where electricity prices vary over the network nodes. No heuristic approach is proposed.

Different combinations of available MLRs are considered in Energy-Aware Routing and Wavelength Assignment (EA-RWA) considered in [121] additionally to the elastic networks. Power consumption of network elements installed along physical links of the network is used as a link weight to select the most energy-efficient lightpath from a set of K-Shortest Paths (KSP).

Weighted Power-Aware Lightpath Routing (WPA-LR) [44] is based on a modified version of the KSP algorithm. Scaled power of OLAs is used as a weight of fiber links. Scaling factor equal to 1 is used for fiber links that are not in use. The selected K paths are evaluated according to the wavelength continuity constraint, and the one with lowest cost is selected for the requested lightpath. If no path satisfies the wavelength continuity constraint, the lightpath request is dropped.

Differentiation of primary and secondary (protection) paths is the basis of the approaches proposed in [45]. The approaches also use the KSP algorithm to precompute available paths between all node pairs. The three considered approaches (Energy-Aware Dedicated Path Protection with Differentiation of primary and secondary paths (EA-DPP-Dif), Energy-Aware Dedicated Path Protection with Mixing Secondary with primary paths (EA-DPP-MixS), and Energy-Aware Dedicated Path Protection (EA-DPP)) use different weights of physical links differentiating sets of links used only by primary paths, links used only by secondary paths, links used by both primary and secondary paths and unused links.

Morea and Perelló propose a MILP with a Connectivity Graph in [3, 48, 49]. The Connectivity Graph is defined as a set of nodes and a set of logical links. The definition of a logical link in [3, 48, 49] is different than the definition in this thesis, since lightpaths are routed over the logical links in [3, 48, 49] (see the variable $x_{de'}^k$ in [3]). No heuristic approach is proposed.

Table 3.4: Energy-Aware Adaptive Routing Solutions (EA-ARSs) targeting power saving in WDM layer.

Method's Name	Targeted Devices	Dynamics of routing	Routing IP	Routing WDM	Physical layer constraints	Constraints on installed devices	Impact on QoS	Computation time	Triggering events	Operation	Network knowl.	Protection consid.	Reconfiguration cost	Future traffic assumption	Control mechanisms
PA-RWA (LBC - LCW) [59, 102, 103], Coiro	OLAs	WDM	n/a	SPR ⁷	consid. (WA)	included	no ⁸	sub-secs.	timer	distr.	global	no	n/c	unknown	n/c
Iterative switching fibers off [92, 93], Coiro	amplifiers ⁹ , switching matrix ¹⁰	WDM	n/a	SPR ¹¹	consid. (WA)	included	n/c	sub-secs.?	traffic increase	distr.	global	SPP	n/c	known	n/c
RWA-Bill ILP [50, 51], Cavdar	OLAs	WDM	n/a	MPR	consid. (WA)	included	n/c	hours	timer (every hour)	centr.	global	no	n/c	known	GMPLS extension
EA-RWA with SLRs and MLRs [121], Vizcaíno	TSPs, OLAs	WDM	n/a	one of KSP	consid. (WA, lightpath length)	included (all devices)	n/c	secs.	new IP traffic demand and its termination	centr.	global	no	n/c	unknown	n/c
WPA-LR [44], Wiatr	OLAs, transceivers, OXCs	WDM	n/a	one of 3-shortest paths ¹²	consid. (WA)	included	wavelength usage	n/r	arrival & termination of lightpath requests	distr.	global	no	n/c	unknown	n/c
EA-DPP-Dif, EA-DPP-MixS, EA-DPP ¹³ [45], Jirattigalachote	OXCs, TXs, RXs, OLAs	WDM	n/a	SPR ¹⁴	consid. (WA)	included	no ⁸	n/r	arrival of lightpath request	centr.	global	DPP with disjoint paths	n/c	unknown	n/c
MILP with Connectivity Graph [3, 48, 49], Morea & Perelló	TSPs, REGs, OLAs	WDM	n/a	MPR	consid. (WA, lightpath length)	included	wavelength usage	n/r	arrival & termination of lightpath requests	distr.	global	Dedicated 1:1	n/c	known	GMPLS extension

⁷LBC with time variant link weights.⁸Blocking probability is analyzed in [102] and [45].⁹In-line, pre- and post-amplifiers are considered in [92, 93].¹⁰MEMS-based switching matrix.¹¹Power of OLAs on the link is used as link weight.¹²Each fiber link is assigned a weight dependent on power of OLAs on the link and whether the link is in use [44].¹³Different assignment of weights (costs) to links is used in EA-DPP-Dif, EA-DPP-MixS, EA-DPP [45].¹⁴Links weights (costs) are assigned.

3.2.3 IP layer

There are many EA-ARs targeting devices in the IP layer for energy saving (see Table 3.5).

The authors of [116] investigate a Power-Aware Traffic Engineering (GreenTE), by proposing MILP coupled with methods to reduce its complexity. In particular, the concept of balancing the load in the “greened” network is proposed, and the signaling issues are thoroughly discussed. QoS is investigated using link utilization, delay, queue length (with ns-2 simulations), and number of MultiProtocol Label Switching (MPLS) tunnels. Reconfiguration aspects and the constraints of the WDM layer are not considered.

The authors of [146–148] propose Virtual Line Card Migration (VLCM), which eliminates rerouted traffic at the IP layer. This is possible by virtualizing the IP logical functionalities, which means moving the IP functionalities from a line card to another one, allowing then to put the first line card into idle state. VLCM introduces a new level in the control plane architecture, and the rerouting takes place at the layer 2 (MPLS/Ethernet). An Energy-Aware Router architecture is also presented.

A depth- d algorithm is proposed [63, 65] as part of the project called MiDORi (Multi-(layer, path and resources) Dynamically Optimized Routing). d determines the number of links that are attempted to be switched off. All possible logical topologies with deactivated d links are generated, and the one with the lowest maximum link utilization is chosen. If utilization of any link exceeds 1, d is decremented and another logical topology with inactive links is searched for until success or until d equals 0. The computation time ranges from 0.01 s up to over a day depending on the value of d and number of nodes in the network. E.g., for a network with 20 nodes the calculation time equals ca. 0.8, 3, 90, 1200 s for d equal to 1, 2, 3 and 4, respectively (Fig. 4 of [65]). The reconfiguration cost is not taken into account.

The Beeler’s algorithm [64, 149] is another algorithm originating from the MiDORi project. It is an efficient implementation of the exhaustive search for the most power-efficient logical topology satisfying given traffic demands. Calculation times of less than 0.01 s have been achieved with an efficient implementation on a parallel data flow type reconfigurable processor for networks with 25–30 nodes (Fig. 7(a) of [64]).

The authors of [133] perform a reality check of Energy-Aware Routing on different network scenarios assuming three different energy models (energy-agnostic, idleEnergy and fully proportional). The constraints on maximum load on logical links are incorporated in the formulated MILP, and analyzed in the solutions. The MILP does not take into account reconfiguration costs.

Minimization of power consumed by router line cards and chassis in each of 12 periods over a day is investigated in [107]. The authors analyze the newly turned-on and shut down line cards (and lightpaths), as well as the percentage of traffic rerouted in the Virtual Topology that is produced as outcome of the MILP solution. The reconfiguration cost is however not included in the objective function of the formulated MILP.

Performance of switch-off and switch-on schemes are compared in [108, 109]. The authors propose a network designed for off-peak hour with the possibility to establish dynamic optical circuits in the peak hour. Then, routing parameters are introduced to trade between energy consumption and route changes. Solutions of the formulated MILPs are presented, without investigating practical heuristics.

Least Flow Algorithm (LFA), Most Power Algorithm (MPA) and the Random algorithm described in [28, 60, 61, 137] attempt to power off entire logical links, with the corresponding line

cards, and possibly also routers, checking if the network is still connected and all IP traffic demands are satisfied according to the updated IP routing. The only difference between the algorithms is the order in which resources are considered, either according to increasing load (LFA), power (MPA) or randomly (Random algorithm). Extension of LFA has been proposed in [112] to switch off single lightpaths instead of the entire logical links.

The ordering issue has been improved in L-Game [113]. The logical links are ordered according to a measure (Shapley value) that defines how much a link is critical in the network, meaning how much traffic cannot be routed if this link is not active.

The previous algorithms do not take into account any cost related to the reconfiguration of the network. Instead, the GA¹⁵ described in [137] searches through random selection, crossover and mutation for a feasible logical topology (which acts as an individual) that reduces a weighted sum of power consumption and reconfiguration costs measured as the amount of IP traffic that has to be rerouted between two consecutive time periods.

A distributed algorithm GReen Distributed Algorithm (GRiDA) is proposed in [150, 151]. It utilizes periodic Link State Advertisements (LSAs) describing configuration and load of links, and broadcasting critical states. The reconfiguration cost in terms of rerouted is not considered in GRiDA itself, but the number of reconfigurations per node ($ON \Rightarrow OFF$ and $OFF \Rightarrow ON$) is analyzed in postprocessing in [150, 151].

Distributed versions of LFA and MPA are proposed in [62]. In the Distributed Least Flow (DLF) and Distributed Most Power (DMP) algorithms each network node maintains three First In First Out (FIFO) queues to store the last links that (i) entered into sleep mode but no LSA confirmed yet that constraints are not violated (*toBeVerified* list); (ii) are in sleep mode and caused no constraint violations (*sleepLinks* list); (iii) caused a violation and thus should not enter sleep mode anymore (*tabu* list with limited length). DLF aims at selecting the least loaded link to be switched off, while DMP aims at selecting the most power-hungry link to be switched off. Similarly to [150, 151], the number of reconfigurations per node is analyzed in [62].

Distributed and Adaptive Interface Switch off for Internet Energy Saving (DAISIES) [36, 59, 94] is another solution that exploits the control mechanisms provided by MPLS to switch off router line cards. The actual amount of traffic carried by each Label Switched Path (LSP) is monitored by the ingress node on a fine granular observation period (e.g., tens of seconds). Then, whenever it goes beyond a prefixed threshold, the ingress node recomputes the path of the LSP and reroutes it updating both the path and the reserved bandwidth (make-before-break). The information about available (unreserved) bandwidth advertised by the Traffic Engineering (TE) routing protocol is used by DAISIES to properly compute link weights, and in turn aggregate the traffic on a reduced set of links.

Energy-Aware Routing (EAR) algorithm [84, 85, 89] considers three types of routers in the network: exporters, importers, and neutral routers. An importer router computes its routing path tree taking into account the Shortest Path Tree (SPT) of the exporter router it is associated with. The three main phases of the EAR algorithm are: (1) exporter router selection; (2) Modified Path Tree selection by an importer router, and determination of the links to be switched off; (3) routing path optimization on the residual topology obtained after the first two steps. There are different exportations which result in putting a link into sleep mode. A generic exportation is referred to as a move in [84, 85, 89]. A set of rules determining whether two moves can be performed one

¹⁵Extension of the GA described in [38].

after another is proposed in [85]. The problem of finding a set of compatible moves minimizing the network power consumption is formulated as a maximum clique problem in a graph. Three heuristics to solve this problem are proposed, i.e., Max_Compatibility [85, 89], Min_Used_Links [85], and QoS-aware_Max_Compatibility [89].

The algorithm Energy Saving based on Algebraic CONnectivity (ESACON) [86] uses algebraic connectivity to choose IP links to be switched off and keeping algebraic connectivity over a given threshold γ_{th} . In the first step, the network links are ordered according to their impact on the algebraic connectivity. In the second step, the links to be switched off are identified starting from the ones that have the lowest impact on the algebraic connectivity.

Energy Saving based on Occurrence of Links (ESOL) leverages the number of occurrences of network links selected by the Open Shortest Path First (OSPF) (Dijkstra) routing [87]. Four different versions of ESOL are proposed in [87], i.e., basic-ESOL, fast-ESOL, $(f + b)$ -ESOL and $(f \times 2)$ -ESOL in order to find a trade-off between computation time and power saving. Basic-ESOL sorts unidirectional links of the logical topology in decreasing order of their occurrences in network paths. Basic-ESOL iteratively selects from this sorted list links that can be switched off until the network remains connected. Fast-ESOL tries to find the largest set of links that can be powered off in one iteration without disconnecting the network. $(f + b)$ -ESOL uses fast-ESOL and basic-ESOL in cascade, so that the output of fast-ESOL is an input to basic-ESOL. $(f \times 2)$ -ESOL repeats f-ESOL in cascade two times.

Energy Saving based on TOPology control (ESTOP) [88] does not use any traffic information in its computation. It is solely based on the topology information and utilizes the notions of edge “betweenness” and graph connectivity to put unnecessary IP links to sleep mode.

A Routing Standby Algorithm [91] aims at putting routers into standby mode. It uses the Floyd-Warshall algorithm (a classical algorithm to find shortest path in a weighted graph), which is stopped before all the nodes are considered. A Demand Reassignment problem (ILP) is solved after stopping the Floyd-Warshall algorithm. Three different versions of the Routing Standby Algorithm are considered, where different sorting criteria (degree-based, betweenness-based, and traffic-based) are used for the initial list of nodes provided to the Floyd-Warshall algorithm.

The idea of reducing energy spent on routing within line cards lies behind the Table Lookup Bypass (TLB) proposed in [37] and extended in [122]. The forwarding engine of some line cards is frozen and incoming packets are forwarded towards a single pre-determined outgoing interface. Table Lookup Capability Bypass Transformation of the classical ILP problem is used to determine incoming, outgoing ports as well as internal arcs and external links of the nodes in the network, which in turn are used to formulate an extended ILP problem. The extended ILP problem determines states of all line cards in the network being either OFF, TLB or FULL. Evaluation of TLB is restricted to a six-node network due to the problem complexity. Larger networks are evaluated using a genetic algorithm proposed in [122]. A weighted sum of power consumed by line cards (in FULL and TLB states) and the average available link capacity is used as a fitness function in [122].

The authors of [66] consider protection in the context of EA-ARSs. A MILP problem covering all considered time intervals is formulated. A heuristic called Energy and Survivability Aware Single Time-period Heuristic (ESA-STH) is also proposed. It separately solves the MILP problem for each single time interval, and not for all considered time intervals at once. A constraint on card reliability is ensured, so that each line card can be switched on (and off) only a limited number of times.

Table 3.5: Energy-Aware Adaptive Routing Solutions (EA-ARSs) targeting power saving in IP layer.

Method's Name	Targeted Devices	Dynamics of routing	Routing IP	Routing WDM	Physical layer constraints	Constraints on installed devices	Impact on QoS	Computation time	Triggering events	Operation	Network knowl.	Protection consid.	Reconfiguration cost	Future traffic assumption	Control mechanisms
GreenTE [116], Zhang	LCs with ports	IP	MPR or SPR ¹⁶	n/c	n/c	included (LCs)	no ¹⁷	65 s – 28 h ¹⁸	timer (5, 15 min)	centr.	global	no	n/c	known	MPLS ext.
VLCM [146–148], Bolla	LCs & TSPs (IP links)	no ¹⁹	SPR, ShPR	n/c	n/c	included (routers and LCs)	IP transparency	secs.	change of IP traffic demand	centr.	global	no	n/c	known	n/c
Depth-d search algorithm [63, 65], Yonezu	IP links	IP	SPR, ShPR, Dijkstra	n/c	n/c	included (IP links)	overprov.	sub-secs.–mins.	n/r	centr.	global	no	n/c	known	n/c
Beeler's algorithm [64, 149], Yamanaka & Takeshita	NICs & switch fabrics	IP	SPR, ShPR, Dijkstra	n/c	n/c	included (IP links)	overprov. & max. IP hop	sub-secs.–secs.	n/r	centr.	global	Disjoint Multi-route divergence	n/c	known	GMPLS ext.
MILP [133], Bianzino	IP links & nodes	IP	MPR	n/c	n/c	included (links & nodes)	overprov.	n/r	change of TM	centr.	global	n/c	n/c	known	n/c
Switch-On & Switch-Off MILPs [108, 109], Caria	LCs & IP ports	IP with changes of LT	SPR	precomputed	n/c	max. number of ports ²⁰	overprov.	n/r	exceeding util. thresholds	centr.	global	no	routing penalty	known	n/c
LFA, MPA, Random [28, 60, 61, 137], Chiaraviglio ²¹	routers, LCs & TSPs (IP links)	IP	SPR, ShPR	n/c	n/c	included (all devices)	overprov.	secs.	change of TM	centr.	global	no	n/c	known	n/c
L-Game [113], Bianzino	LCs & TSPs (IP links)	IP	ShPR	n/c	n/c	included (all devices)	Shapley value	secs.	change of TM	centr.	global	no	n/c	known	n/c
GA Bonetto [137]	LCs & TSPs (lightpaths)	IP	SPR, ShPR	n/c	n/c	included (routers and LCs)	overprov.	mins.	change of TM	centr.	global	no	rerouted traffic	known	n/c

Continued on next page

¹⁶One of KSP in simplified MILP.¹⁷Routing stability is analyzed in [116].¹⁸Computation times are provided for the AT&T network with 115 nodes and 296 links in [116].¹⁹Rerouting takes place in an intermediate layer between IP and WDM in [148].²⁰No capacity constraints in the transport network in [108, 109].²¹Note that the only difference between LFA, MPA, Random is the policy for sorting nodes and links [61].

Table 3.5 Energy-Aware Adaptive Routing Solutions (EA-ARSs) targeting power saving in IP layer – continued from previous page.

Method's Name	Targeted Devices	Dynamics of routing	Routing IP	Routing WDM	Physical layer constraints	Constraints on installed devices	Impact on QoS	Computation time	Triggering events	Operation	Network knowl.	Protection consid.	Reconfi- guration cost	Future traffic assumption	Control mech- anisms
GRiDA [150, 151], Bianzino	router ports & elec- tronic REGs (lightpaths)	IP	SPR, ShPR, BFS	n/c	n/c	included	over- prov.	n/r ²²	LSA arrival	distr.	global (LSAs)	no	no ²³	un- known	signal. with LSAs of OSPF or ISIS
DLF & DMP [62], Bianzino	NICs and REGs (IP links)	IP	ShPR	n/c	n/c	included (all de- vices)	over- prov., connect. check	n/r	every 10– 20 s	distr.	local	no	n/c ²³	un- known	periodic LSAs
DAISIES [36, 59, 94], Coiro	LCs & TSPs (lightpaths)	IP ²⁴	SPR, ShPR, Dijkstra	n/c ²⁵	n/c ²⁶	included (routers and LCs)	over- prov.	sub- secs.	change of IP traffic demand	distr.	global	no	n/c	un- known	MPLS rerouting
EAR ²⁷ [84, 85, 89], Cianfrani	NICs	IP	SPR, ShPR, Dijkstra	n/c	n/c	included (IP links)	over- prov.	n/r ²⁸	exceed- ing util. thresh- olds	distr.	global	no	traffic vari- ation on a link due to exportation	un- known	LSAs of the OSPF
ESACON [86], Cuomo	LCs	IP	SPR?, ShPR, Dijkstra	n/c	n/c	included (IP links)	alge- braic connec- tivity thresh- old γ_{th}	n/r ²⁹	n/c	distr./ centr.	global	no	n/c	un- known	OSPF control messages
ESOL ³⁰ [87], Cuomo	capacity modules (2.5 Gbps)	IP	SPR, ShPR, Dijkstra	n/c	n/c	included (IP links)	n/c	n/r ³¹	n/c	distr.	global	no	n/c	known	LSAs?

Continued on next page

²²The time complexity of the solution scales linearly with number of nodes and exponentially with node degree [151].

²³Number of OFF \Rightarrow ON and ON \Rightarrow OFF events is analyzed in [62, 150, 151].

²⁴Possibility of changing the logical topology is considered in [94], but realization of the links (including routing of lightpaths) is not considered.

²⁵LBC was applied to the DAISIES solution in [59].

²⁶LCW was applied to the DAISIES solution in [59].

²⁷The term Energy Saving IP Routing (ESIR) instead of EAR is used in [89]. Max_Compatibility heuristic is considered in [85, 89]. QoS-aware_Max_Compatibility heuristic is considered in [89].

²⁸No increase of complexity with respect to OSPF is reported in [84].

²⁹Complexity of $\mathcal{O}(E \cdot N^2)$ is reported in [86], where E is the number of bidirectional IP links, and N is the number of routers in the network.

³⁰Four versions of the ESOL algorithm are proposed in [87], i.e., basic-ESOL, fast-ESOL, $(f + b)$ -ESOL and $(f \times 2)$ -ESOL.

³¹Complexity of $\mathcal{O}(Niter \cdot N^2 \log_2 N)$ is reported in [87], where $Niter$ is the number of iterations, and N is the number of routers in the network.

Table 3.5 Energy-Aware Adaptive Routing Solutions (EA-ARs) targeting power saving in IP layer – continued from previous page.

Method's Name	Targeted Devices	Dynamics of routing	Routing IP	Routing WDM	Physical layer constraints	Constraints on installed devices	Impact on QoS	Computation time	Triggering events	Operation	Network knowl.	Protection consid.	Reconfiguration cost	Future traffic assumption	Control mechanisms
ESTOP Cuomo [88],	capacity modules (2.5 Gbps)	IP	SPR, ShPR, Dijkstra	n/c	n/c	included (IP links)	n/c ³²	n/r ³³	n/c	distr.	global	quick activation of links	n/c	unknown	LSAs of the OSPF, make-before-break
Routing Standby Algorithm ³⁴ [91], Cianfrani	routers	IP	SPR?, ShPR, Floyd-Warshall	n/c	n/c	included (IP links)	min. of load in Demand Reassignment	n/r	n/c	centr.	global	quick activation of routers (standby)	n/c	known	n/c
TLB ³⁵ [37, 122], Coiro	LCs (off or TLB state)	IP	SPR	n/c	n/c	included (IP links)	overprov. ³⁶	n/r ³⁷	n/c	centr.	global	TLB state	n/c	known	n/c
ESA-STH (MILP) [66] (Addis, sustainIT 2012)	LCs & chassis & TSPs (lightpaths)	IP	SPR	n/c	n/c	included (routers and LCs)	overprov.	n/r	change of TM	centr.	global	IP, shared	on/off limit	known	n/c

³²Increase of path length in terms of IP hops, mean utilization of the links, and fairness of traffic distribution on all links are analyzed in postprocessing in [88].

³³Complexity of $\mathcal{O}(M \cdot N^2)$ is reported in [88], where M is the number of iterations, and N is the number of routers in the network.

³⁴Three versions of the Routing Standby Algorithm are considered in [91], where different sorting criteria (degree-based, betweenness-based, and traffic-based) are used for the initial list of nodes provided to the Floyd-Warshall algorithm.

³⁵A MILP formulation [37, 122] and a heuristic based on a genetic algorithm [122] are proposed. Both use the TLB.

³⁶The average path length in terms of IP hops is analyzed in postprocessing in [37].

³⁷Complexity of $\mathcal{O}(|\nu|^5)$ is reported in [37], where $|\nu|$ is the number of nodes in the network.

3.2.4 IP and WDM layers

Covering both the IP and WDM layers is challenging for EA-ARSSs. Summary of the approaches targeting energy saving in both layers is presented in Table 3.6.

Network Optimization Algorithm is reported in [82]. This centralized solution validates physical layer constraints (wavelength assignment and lightpath length). However, a detailed description of the algorithm (e.g., with a pseudocode or a flowchart) is missing.

Auxiliary graph heuristics Request Size Based (RSB) and Link Utilization Based (LUB), Red-Demand-Based (RDB) and Total-Demand-Based (TDB) are proposed in [56–58]. The auxiliary graph consists of access layer, lightpath layer, chassis layer, module layer, port layer and a layer for each wavelength. Link weights in the connectivity graph are assigned values proportional to power consumption of corresponding devices. The RSB and LUB heuristics [57, 58] share the same concept, where node pairs are sorted first (according to total demand for RSB or link utilization for LUB). The traffic demands corresponding to the sorted node pairs are routed over the shortest path in the auxiliary graph, where lightpaths of capacity lower than the traffic demands are removed. Differentiated energy saving traffic grooming is considered in [56]. Two different classes of traffic are considered (Red and Green). The RDB heuristic runs the RSB algorithm for the Red traffic demands first, and then runs RSB again for the Green traffic demands. In the TDB heuristic, all traffic demands (Red and Green) are sorted on the same node pair list according to the total demand (Red + Green). The traffic demands are served according to the list, with the Red traffic demands having priority to the Green traffic demands within the same node pair.

The next two approaches also utilize connectivity graphs, however the problems are formulated as MILP problems. The focus of [55] is set on reach related power consumption of 100G transponder. The formulation in [52–54, 106] uses the network design MILP formulation reported in Section 3.1.2 with additional constraints on the installed devices in the network.

Eventually, both the WDM and IP layers are covered in [59], where two steps are considered. First, the IP routing and logical topology are determined using DAISIES (see Section 3.2.3). Second, PA-RWA (LBC and LCW, see Section 3.2.2) is used to find realization of the logical topology from the first step.

Table 3.6: EA-ARs targeting power saving in IP and WDM layers.

Method's Name	Targeted Devices	Dynamics of routing	Routing IP	Routing WDM	Physical layer constraints	Constraints on installed devices	Impact on QoS	Computation time	Triggering events	Operation	Network knowl.	Protection consid.	Reconfiguration cost	Future traffic assumption	Control mechanisms
Network Optimization Algorithm [82], Klekamp	IP routers, LCs, TSPs, OLAs, WDM shelves, filters AWG	IP & WDM	SPR?	SPR, ShPR	consid. (WA, lightpath length)	n/r	n/c	n/r	n/r	centr.	global	no	n/c	known	n/c
Auxiliary Graph heuristics RSB and LUB, RDB and TDB [56–58], Hasan & Farahmand	router chassis, ports	IP & WDM	SPR in Auxiliary Graph	SPR in Auxiliary Graph	consid. (WA)	included (all devices)	differentiated traffic	n/r	timer (every 30 min. [58])	centr.	global	no	n/c	known	n/c
MILP with Connectivity Graph [55], Rizzelli	TSPs, REGs, IP processing	IP & WDM	MPR	MPR	consid. (WA, lightpath length)	included (TSPs and REGs)	n/c	hours	total traffic increase by 5%	centr.	global	no	n/c	known	Cited [21] – GMPLS ext.
MILP with Connectivity Graph [52–54, 106], Rizzelli & Morea	TSPs, REGs, routers ³⁸ , interconnecting fabrics, LCs	IP & WDM	MPR	MPR	consid. (WA, lightpath length)	included (all devices)	no ³⁹	hours	total traffic increase by 5%	centr.	global	no	n/c	known	n/c
MILP [107, 127], Zhang	chassis, LCs & OXCs	IP & WDM	MPR	MPR	consid.	included ⁴⁰	n/c	hours	change of a TM	centr.	global	n/c	top-down scheme	known	n/c
Two steps: DAISIES + PA-RWA (LBC and LCW) [59], Coiro	LCs, switching fabrics, OLAs	IP & WDM	ShPR ⁴¹	ShPR ⁴²	consid. (WA)	included (all devices)	n/c	sub-secs.	change of IP traffic demand	distr.	global	no	n/c	known	n/c

³⁸Slave shelves, slave racks, router chassis are distinguished in [52–54, 106].

³⁹Maximum utilization of resources is analyzed in [106]

⁴⁰The mathematical formulation in [107] ensures the constraint on the installed fibers, but only partly addresses the constraint on the installed chassis, LCs and OXCs.

⁴¹Time variant link weights at the IP layer are used in [59].

⁴²LBC with time variant link weights at the WDM layer are used in [59].

3.2.5 Concluding remarks and thesis contributions

The highest number of EA-ARSs target the IP devices for power saving. There is only one EA-ARS approach covering both the IP and WDM layers [57, 58]. It uses the notion of connectivity graph. Connectivity graphs are used also in the MILPs [55] and [52–54, 106] by Rizzelli et al. However MILP is not suitable in the dynamic network operation targeted by EA-ARSs. Relaxation of the problem by splitting it into two (algorithmic) steps is addressed in [59]. The Network Optimization Algorithm reported in [82] is promising, however its details are missing in [82]. Almost all the works target line cards to be on/off switched (or put into standby mode). Routers and OXC are rarely addressed, since possibility of their dynamic operation is questionable. OLAs are the most frequently targeted devices in the WDM layer. Heuristics use mostly SPR in the IP layer. Lightpaths between the same node pair are usually allowed to traverse different paths (MPR in the WDM layer). Physical layer constraints are followed to various extents. Wavelength assignment is the most frequently addressed constraint. Some works consider also the constraint on maximum lightpath length. Due to classification of the related work between EA-ND and EA-ARSs, all considered EA-ARSs obey the constraint on the installed devices. However, the constraints are not always applied to all types of devices. Similarly to the EA-ND approaches, overprovisioning is the most frequently used method for ensuring QoS. There are however also other methods such as constraints on connectivity and the use of differentiated traffic. The reported computation times are only indicative. However, some heuristics obtain solutions within the time scale of subseconds–seconds even if their implementation is not optimal. Triggering events are mostly related to the change of traffic (either the whole traffic matrix or load on the lightpaths). In some works new solutions are computed periodically, what corresponds to an expiry of a timer. There are only few approaches which can be operated in a distributed way. Most of the approaches require global network knowledge.

Protection is considered by a few EA-ARSs, however there is still space for research in this area particularly in the class of EA-ARSs targeting energy saving in the IP layer. Several promising EA-ARSs (e.g., DAISIES [36, 59, 94]) do not address the protection at all, and a simplistic assumption of quick activation of links and routers is taken in some other works [88, 91],

Several related works analyze the influence of application of EA-ARSs using as evaluation metrics e.g., blocking rate of lightpath requests or average hop count that a traffic demand traverses from its source to its target. However, only few EA-ARSs consider reconfiguration cost during their operation (routing penalty [108, 109], rerouted traffic [84, 137], on/off limit [66], a top-down approach [107]).

EA-ARSs use traffic as input data. The traffic is usually assumed to be known in the related works. The assumption of unknown traffic is more frequently made by the EA-ARSs targeting WDM devices to be switched off. This is caused by the fact that traffic in the WDM layer is usually modeled as arriving lightpath requests with given holding time, and not as a traffic matrix of certain time granularity. A few works address the control mechanisms (more details in Section 6.2). Evaluation of the EA-ARSs together with appropriate control mechanisms has not been well investigated yet.

The contribution of this thesis are the following. First, we evaluate the potential of power saving by the EA-ARSs in the IP layer and consider different levels of freedom for rerouting of traffic demands and lightpaths (Chapter 4). The approach when no rerouting both in the IP and WDM layer is allowed is a simple distributed EA-ARS which requires only local knowledge for its operation.

Second, we propose an adaptive EA-ARS, where network configuration from previous time period is used to determine the network configuration for the following time period (Chapter 5). Clear triggering events are described, which help timely reconfigure the network minding QoS. Computation time of the proposed algorithm is reasonable. Evaluation studies are performed using as realistic input data as possible (Chapter 2).

Third, implementation issues for the EA-ARSs are discussed, and two of the proposed energy saving approaches are implemented on a testbed (Chapter 6). We demonstrate that it is feasible to dynamically and remotely reconfigure the network, save energy and avoid traffic loss using off-the-shelf equipment.

4 Potential energy savings in the IP-over-WDM networks

As a first step, we estimate how much energy could be potentially saved in the Internet Protocol (IP)-over-Wavelength Division Multiplexing (WDM) network by switching off line cards in low-demand hours. We first present the formulation of the optimization model to design the Static Base Network (SBN). Next, we present three approaches to saving energy with different levels of freedom of rerouting. Considered scenarios and results follow, and eventually conclusions are drawn. This chapter is based on [136, 152].

4.1 Design of Static Base Network

To the best of our knowledge, there is no complete and publicly available data about the networks today. The complete data includes physical topology (physical nodes and links with their capacity in terms of fibers and number of WDM channels per fiber), the logical topology (logical nodes and links with their capacity, as well as realization including routing of lightpaths over the physical topology), and the routing of IP demands over the logical topology. Information about the network elements installed in the network, as well as their configuration is required in order to evaluate potential energy savings.

We formulate the network design problem as a Mixed-Integer Linear Programming (MILP) in order to determine the required information for a network of today. Cost minimization is set as the objective of the MILP. Capital Expenditures (CapEx) or power consumption can be considered as cost. We use CapEx minimization as the objective for the design of the network of today, which is static and therefore referred to as the SBN.

Designing a cost-minimal multi-layer network that allows to realize a given traffic matrix is a highly complex problem which is far from being solved yet, see for instance [153–157]. The problem becomes even harder if multiple traffic matrices are to be considered, leading to robust multi-layer network design problems [158, 159]. Most attempts to solve robust network dimensioning problems, however, assume single-layer networks, see [160–162] and references therein.

Given detailed traffic measurements, our approach is based on constructing a single traffic matrix that refers to all peak demands over time. The SBN is then dimensioned with respect to this maximum matrix which ensures that every single traffic scenario can be realized. The SBN can be considered being cost-minimal among all networks that allow to route the constructed maximum traffic matrix. Note that although common in practice, this approach is potentially producing over-provisioned networks. There might be cheaper topologies that cannot accommodate the maximum traffic matrix but all single traffic scenarios (with static IP routing). It is also typically unlikely to have all demands at their peak simultaneously. For our study, however, this approach based on a maximized matrix is reasonable since our main goal is to compare the three different energy saving

approaches among each other rather than to provide the cheapest among all robust SBNs.

Let V be the set of all demand end-nodes¹ and let $d^{ab}(t)$ be the undirected demand value for each pair of nodes $(a, b) \in V \times V, a < b$ and each point in time $t \in T$, where T is the set of considered points in time (time periods). We compute each element of the maximum demand matrix by

$$d_{SBN}^{ab} = \max_{t \in T} d^{ab}(t), \quad (4.1)$$

and calculate a minimum-cost IP-over-WDM network which satisfies this maximum demand matrix.

Our model used to cost-optimally design the SBN is close to the one used in [153, 156, 163–165]. We optimize both network layers in an integrated step. The model comprises all relevant sources of installation cost both in the IP and the WDM layer. Modified versions of this model are later used to evaluate the energy savings in different traffic demand scenarios taking the designed SBN as input (see Section 4.2 and Appendix B).

Parameters Assuming all network elements to be bidirectional, we model the optical layer by an undirected physical supply network (V, E) consisting of the nodes V and the physical supply links E . Every node $i \in V$ can be equipped with an IP router out of the set N of IP routers. Every router $n \in N$ has a maximum switching capacity of \mathcal{R}^n and a cost of \mathcal{C}^n . Every physical supply link $e \in E$ can operate an arbitrary number of fibers at cost \mathcal{C}^e per fiber, each supporting B wavelength channels. For simplicity we do not consider different optical nodes to be available. Instead we assume a pre-installed Optical Cross-Connect (OXC) of infinite capacity at every network node. The model (4.3) can be easily extended by a physical node model though. Also note that we include some cost for OXCs in the cost of fibers, see Section 4.3.

For every node pair $(i, j) \in V \times V, i < j$, the set $P_{(i,j)}$ denotes all admissible routing paths in (V, E) between nodes i and j , which can be used to realize lightpaths. Let P be the union of all these paths and P_i the set of all paths ending at node i . Every path $p \in P$ can be equipped with multiple lightpaths of capacity W . Each bitrate unit W on a path p incurs the cost \mathcal{C}^{LC} of line card interfaces at the end-nodes of p and consumes one wavelength channel in the physical network on every physical link of the path. All lightpaths established between a node pair $(i, j) \in V \times V$ forms a logical link. Set of all logical links is denoted as L , and the corresponding logical topology as (V, L) .

Demands and commodities We assume that IP traffic can be arbitrarily split and routed via multiple logical paths. This is modeled by a standard so-called splittable multi-commodity flow [166] on the IP network layer. We remark that the main motivation behind this assumption is computational tractability since any more restrictive routing assumption (e.g., single-path flow or path length restrictions) would involve using discrete flow variables and/or additional flow constraints. From the practical point of view, we refer the reader to [167] for a discussion on the similarity of load distribution in a network using multi-commodity flows and a network using the Open Shortest Path First (OSPF) protocol with Equal-Cost Multi-Path (ECMP) routing and clever weight setting.

We introduce commodities based on the given point-to-point demands $d_{SBN}^{ab}, (a, b) \in V \times V, a < b$ in order to model a multi-commodity flow. There are mainly two approaches related to the definition of commodities [166, 168–170]. The first is to consider one commodity for every non-zero point-to-point demand. This approach results in so-called disaggregated formulations which can become huge already for small networks. The number of variables and constraints in such

¹We use indices i (and j) to refer to a node (and a node pair) in general, and a and b in the context of traffic demands.

models is in the order of $\mathcal{O}(|V|^4)$ and $\mathcal{O}(|V|^3)$, respectively, just for modeling the flow. For smaller models and to reduce computation times it is common to aggregate demands at common source nodes which leads to commodities having one source but several target nodes. This modeling trick reduces the number of commodities to at most $|V|$ and the number of variables and constraints in the multi-commodity flow model to $\mathcal{O}(|V|^3)$ and $\mathcal{O}(|V|^2)$, respectively.

In the following the set of commodities $K \subseteq V$ corresponds to those nodes in V that are source of at least one demand. For commodity $k \in K$ and every node $i \in V$ the net demand value is defined according to Eq. (4.2), where all demands whose source is $k \in V$ are subsumed.

$$d_i^k = \begin{cases} \sum_{b \in V} d_{SBN}^{ib} & \text{for } i = k \\ -d_{SBN}^{ki} & \text{otherwise} \end{cases} \quad (4.2)$$

It holds that $\sum_{i \in V} d_i^k = 0$ for all $k \in K$. The total demand value d_i^{SBN} of a network node i is given by the sum of all demands having its source or target in i , that is, $d_i^{SBN} = \sum_{k \in K} |d_i^k|$.

Variables The flow variables $f_{ij}^k, f_{ji}^k \in \mathbb{R}_+$ describe the directed flow for commodity k on the logical link between i and j in both directions. Notice that the flow variables are not defined for individual lightpaths. The variables aggregate the IP traffic on all lightpaths with end-nodes i and j . This is possible because the actual physical representation of a logical link does not matter for the IP routing. Only the total capacity between any two nodes is of interest for an IP demand in our model. Also notice that by the definition of the commodities above the flow variables aggregate IP traffic with the same source node. The distribution of logical link flow to the chosen physical representations and also the disaggregation of the commodity flows to individual demand flows can be done in a postprocessing step, as explained below and in [156, 163]. Both aggregation techniques significantly reduce the size of the model compared to considering flow variables on individual physical representations and for individual point-to-point demands.

Variables $y_p^{SBN} \in \mathbb{Z}_+$ count the number of lightpaths realized on $p \in P$. Similarly, $y_e^{SBN} \in \mathbb{Z}_+$ denotes the number of fibers installed on physical supply link $e \in E$. The binary variable $x_i^n \in \{0, 1\}$ states whether or not router $n \in N$ is installed at node $i \in V$.

Model The problem of minimizing the cost for a feasible network configuration and routing satisfying the demand matrix D_{SBN} can be formulated as the MILP (4.3).

$$\min \sum_{i \in V, n \in N} C^n x_i^n + C^{LC} \sum_{p \in P} y_p^{SBN} + \sum_{e \in E} C^e y_e^{SBN} \quad (4.3a)$$

$$\sum_{j \in V \setminus \{i\}} (f_{ij}^k - f_{ji}^k) = d_i^k, \quad i \in V, k \in K \quad (4.3b)$$

$$\sum_{p \in P(i,j)} W y_p^{SBN} - \sum_{k \in K} (f_{ij}^k + f_{ji}^k) \geq 0, \quad (i, j) \in V \times V \quad (4.3c)$$

$$\sum_{n \in N} \mathcal{R}^n x_i^n - \sum_{p \in P_i} W y_p^{SBN} \geq d_i^{SBN}, \quad i \in V \quad (4.3d)$$

$$\sum_{n \in N} x_i^n \leq 1, \quad i \in V \quad (4.3e)$$

$$B y_e^{SBN} - \sum_{p \in P: e \in p} y_p^{SBN} \geq 0, \quad e \in E \quad (4.3f)$$

$$f_{ij}^k, f_{ji}^k \in \mathbb{R}_+, \quad y_p^{SBN}, y_e^{SBN} \in \mathbb{Z}_+, \quad x_i^n \in \{0, 1\} \quad (4.3g)$$

The objective (4.3a) is to minimize the cost of installed routers, line cards and fibers in the network. Equations (4.3b) are the flow conservation constraints for every node and commodity, formulated on the complete logical layer $V \times V$. Inequalities (4.3c) choose a subset of paths between the nodes i and j and install enough capacity to accommodate all the logical link flow corresponding to (i, j) . The logical node capacity constraints (4.3d) make sure that the capacity of a node suffices to switch all the incoming traffic, including the emanating demand. The constraints (4.3e) select one router configuration at every node. Eventually, the physical link capacity constraints (4.3f) make sure that the number of available wavelengths on a fiber is not exceeded.

The flow for commodities $k \in K$ on logical link $(i, j) \in V \times V$ has to be disaggregated in two steps in postprocessing. First, we have to disaggregate the flow from the logical links (i, j) to all physical representations $p \in P$ with end-nodes i and j (in both directions). By (4.3c) this can be done respecting the installed capacities $W \cdot y_p^{SBN}$ for $e \in P_{(i,j)}$. Let f_p^k denote the flow on path p for commodity $k \in K$. In a second step, the commodity flow f_p^k has to be disaggregated to a flow for every individual point-to-point demand $(a, b) \in V \times V$, where a corresponds to the source of commodity k and b to one of its targets. Since commodities have been aggregated from demands at common source nodes, this can be done in a greedy manner by iteratively subtracting the necessary individual demand flow for every target b from the given aggregated commodity flow. We denote by f_p^{ab} the flow for demand (a, b) with value d_{SBN}^{ab} on path p .

4.2 Approaches to save energy

Given the SBN designed with the model from Section 4.1, we propose three (classes of) approaches to save energy in low-demand hours.

Fixed Upper Fixed Lower (FUFL): Both the routing of IP traffic in the upper logical layer and the realization of lightpaths in the lower WDM layer are fixed over time. Demands are routed as in the SBN using the same lightpaths with the same percentual splitting as in the SBN. We allow to shift traffic between parallel lightpaths though. Line cards of empty lightpaths are switched off.

Dynamic Upper Fixed Lower (DUFL): The logical topology (including the realization of lightpaths) is fixed as in FUFL (Fixed Lower), but the routing of IP traffic can be changed (Dynamic Upper). In every demand scenario, we aim at routing the IP demands in the logical topology in such a way that as many lightpaths as possible are emptied in order to switch off the corresponding line cards.

Dynamic Upper Dynamic Lower (DUDL): Both the routing of the IP traffic in the logical layer and the realization of lightpaths in the physical layer can be changed over time (including adding lightpaths not existing in the SBN), with the restriction that the number of installed line cards at each IP router must not be exceeded. The number of used line cards is minimized by jointly optimizing the routing in the IP and WDM layers.

A Fixed Upper Dynamic Lower approach is not feasible, since the IP routing has to react to changes of the logical topology. Note that the terms Fixed Lower and Dynamic Lower apply to dynamics of the realization of the logical topology. Idle lightpaths are dynamically switched off in all the considered approaches.

Fig. 4.1 shows a simple example illustrating how the approaches FUFL, DUFL and DUDL can decrease the number of active line cards in a low-demand hour. We limit the number of nodes, links and demands in this example in order to keep it clear and easy to follow. The physical fiber

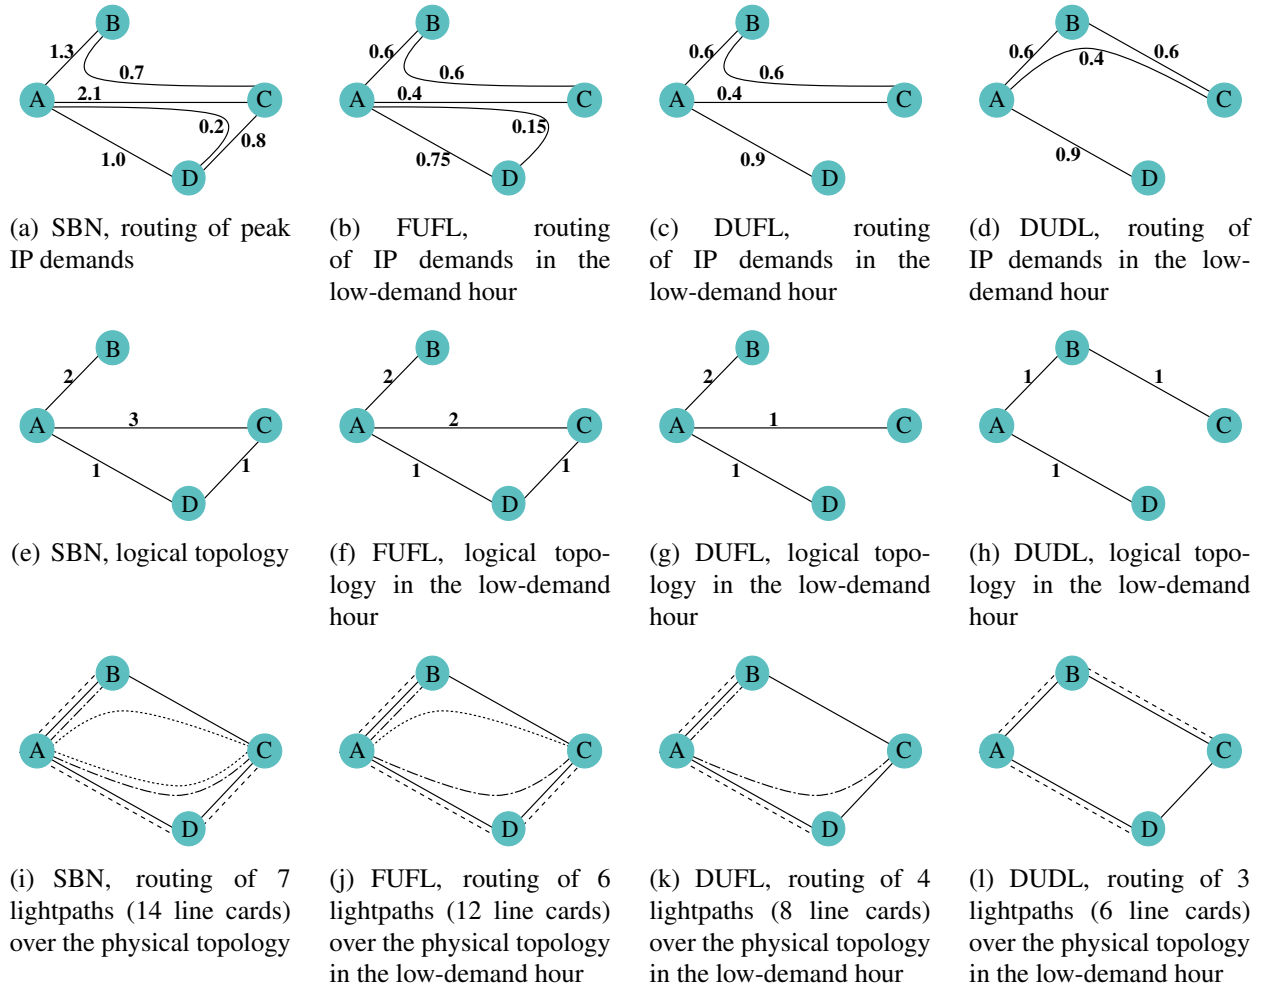


Figure 4.1: A simple example showing how the approaches FUFL, DUFL and DUDL decrease the number of active line cards. There are five peak demands ($AB = 1.3$, $AC = 2.1$, $AD = 1.2$, $BC = 0.7$, $CD = 0.8$), which decrease in the low-demand hour ($AB = 0.6$, $AC = 0.4$, $AD = 0.9$, $BC = 0.6$, $CD = 0.0$). The physical topology is fixed (solid lines in subfigures (i)–(l), each line corresponds to a single fiber of capacity $B = 3$ wavelengths). The granularity of the logical link capacity is $W = 1$ (normalized capacity, two line cards). In the low-demand hour 2, 6, and 8 line cards are saved with FUFL, DUFL, and DUDL, respectively.

installation and the hardware configurations at the nodes from the SBN are fixed for all approaches. New line cards must not be installed. Two traffic matrices are considered. The SBN is designed to accommodate the peak traffic matrix. The approaches FUFL, DUFL and DUDL are applied in the low-demand hour starting from the SBN. The columns of the subfigures in Fig. 4.1 correspond to the SBN, FUFL, DUFL and DUDL respectively. The first row shows the routing of the IP demands, the second one depicts the logical topologies, and the last one illustrates the routing of lightpaths over the physical topology. FUFL can switch off a lightpath between nodes A and C (the dotted one routed via D in Fig. 4.1(i)) due to a decrease of traffic. DUFL changes the routing of IP demands to additionally make the lightpaths AC (dotted, routed via B in Fig. 4.1(i)) and CD idle. Eventually, DUDL adds a logical link between nodes B and C , which does not exist in the SBN (compare Fig. 4.1(e) and (h)). This additional link and the new routing of the IP demands allows to further decrease the number of active line cards.

FUFL is the most restrictive approach. It is the easiest to be realized in practice since it does not require any optimization but only monitoring of the lightpath utilization. Decisions on switching line cards on and off can be taken locally. Its drawback is to rely on the routing defined by the SBN, no matter what it is (single-path, weighted multi-path, etc.). The routing taken from the SBN can be suboptimal especially in low-demand hours. A similar concept to FUFL (extended to different types of networks) was used in [123–126] after the publications [136, 152].

In contrast, DUFL and DUDL with the objective of minimizing number of lightpaths are NP-hard optimization problems, as they generalize the uncapacitated fixed charge flow problem [171, 172]. DUFL is a single-layer network design problem which can be solved to optimality in a reasonable amount of time in practice, see Section 4.4 and [170, 173]. DUDL is a computational challenge since it involves optimizing two coupled network layers simultaneously, similar to designing the SBN, see [153–157].

Dynamics in the IP routing (DUFL) may allow more line cards to be switched off, compared to FUFL, by choosing a smart IP routing in each demand scenario, but it may lead to instabilities of connection-oriented protocols (e.g., due to overtaking of packets upon the change of the IP routing). Moreover, decisions about the IP routing changes need to be forwarded to all involved routers. Even more signaling is needed for additional dynamics in the WDM layer (DUDL). It has to be ensured that no packets are lost in the reconfiguration phase, when lightpaths are torn down. Generalized MultiProtocol Label Switching (GMPLS) including the traffic engineering extensions is a potential candidate for controlling and managing the paths. The reconfiguration itself is non-trivial though since it requires the use of OXCs to dynamically change logical links (typically realized by point-to-point connections nowadays). Energy-efficient network design may have an influence on resilience and Quality of Service (QoS) in the network in terms of packet delay, jitter, packet loss as well as on network throughput, since switching off line cards decreases the capacity of the network. Moreover, network devices being repeatedly switched on and off may be more prone to failures.

More formally, the model (4.3) can be adapted to evaluate the possible energy savings for FUFL, DUFL and DUDL under dynamically changing traffic demands (indicated by t). The complete model variants (DUFL and DUDL) are presented in Appendix B.

For all approaches we fix the physical network by fixing the variables y_e^{SBN} to the values from the SBN. We also fix the variables x_i^n of the installed IP routers together with all installed line cards at each node $i \in V$ according to the SBN solution. The capacity of the IP router installed at node i is denoted by R_i^{SBN} .

FUFL Consider a demand between nodes a and b with base value d_{SBN}^{ab} . In each low-demand hour $t \in T$ where this demand has value $d^{ab}(t)$, we reduce the flow on each path p used to transport this demand by the factor $d^{ab}(t)/d_{SBN}^{ab} \in [0, 1]$ and reduce the capacity on the path accordingly. As all flows for this demand are scaled by the same factor, the relative share of traffic on each used path remains the same as before.

To state it more precisely, consider a path p used to route demand d_{SBN}^{ab} . In scenario $t \in T$, we reduce the flow from f_p^{ab} to

$$f_p^{ab}(t) = f_p^{ab} \cdot d^{ab}(t)/d_{SBN}^{ab}.$$

Let

$$f_p(t) = \sum_{(a,b) \in V \times V} f_p^{ab}(t)$$

be the total flow on path p after reducing the flow for all demands. Then the capacity on that path can be reduced from y_p^{SBN} to

$$y_p(t) = \left\lceil \frac{f_p(t)}{W} \right\rceil,$$

where y_p^{SBN} is the capacity of all lightpaths established on the physical path p in the SBN. The variables $y_p(t)$ count the number of lightpaths on $p \in P$ at time period $t \in T$. Hence this procedure corresponds to shifting traffic between parallel lightpaths, that is, lightpaths with the same realization $p \in P$. No optimization is needed. From a practical perspective alternative FUFL principles might be of interest. Traffic could be shifted between all paths $p \in P$ having the same end-nodes (i, j) ignoring physical representations. This could reduce the number of active lightpaths even more and completely relies on information available at the IP-layer. Additionally adjusting the splitting of demands across IP multi-paths could further contribute to energy savings. For ease of exposition we neglect to study these FUFL variants here.

DUFL For every $t \in T$, we compute new IP routing by using a variant of model (4.3). The logical link (all physical paths on the logical link) capacity variables $y_p(t)$ can be reduced compared to the corresponding ones from SBN (y_p^{SBN}), but not augmented. This is ensured by adding the constraint $y_p(t) \leq y_p^{SBN}$ to (4.3) for all $p \in P$. The IP flow can be rerouted without any further restrictions such as to minimize the number of active line cards. Hence an energy optimal routing for every time period $t \in T$ can be computed by fixing y_e^{SBN} and x_i^n , and bounding $y_p(t)$ in (4.3) as described above, using the demand matrix $D(t)$ instead of D_{SBN} , and changing the objective function to minimize the number of active lightpaths $\sum_{p \in P} y_p(t)$. The detailed model (B.1) to reoptimize the routing in every time step using the strategy DUFL can be found in the Appendix B.

DUDL We are allowed to change the IP routing as well as the logical topology including physical representation of the lightpaths. But we cannot install new line cards at IP routers. Similarly to DUFL, we use a variant of model (4.3) to compute an energy-efficient network in each demand scenario $t \in T$. In contrast to the previous case, the augmentation of logical link capacity variables $y_p(t)$ is allowed in addition to changing the flow variables. In order not to exceed the number of line cards installed at each node in the SBN, we add the constraints

$$\sum_{p \in P_i} y_p(t) \leq \sum_{p \in P_i} y_p^{SBN}$$

for every node $i \in V$. We then minimize the number of used line cards for every demand scenario using the same objective as for DUFL, see model (B.2) in the Appendix B.





Figure 4.2: Physical supply network topology, and source traffic distribution. The area of a node represents its emanating demand. The DFN demands (a) are Frankfurt-centralized in contrast to the DWG demands (b). Notice that the New York node is omitted in the Géant figure (d).

4.3 Scenarios

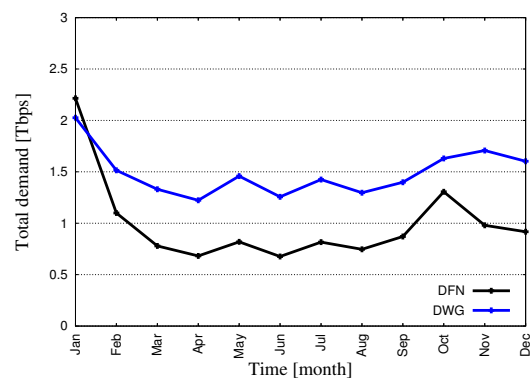
4.3.1 Topologies and traffic demands

We used three physical supply network topologies (Germany17, Abilene, Géant) and a subset of corresponding traffic data [4] described in Section 2.5.

For each single physical supply network topology, we precalculated the set $P_{(i,j)}$ of up to 50 shortest paths for potential lightpaths for every node pair (i, j) . The length of a physical link was computed by using the spherical distance of its end-nodes. The paths were limited in total physical link length to 3000 km. There are three physical links in the Géant network with a real length greater than 3000 km (Israel–Netherlands, New York–Austria and New York–UK). We set these links to a length of exactly 3000 km. There is thus only the direct path possible between these nodes. This way we implicitly assume to provide the necessary regenerators on these three links



(a) Germany17, 1 day over February 2005, 3 Tbps



(b) Germany17, 1 month over 2004, 3 Tbps

Figure 4.3: Total demand over time in Tbps for the scenario with 3 Tbps maximum total demand, every day of February 2005, and every month of 2004. The DWG matrices show the expected behavior over time with peaks during the week, low traffic on weekends (a), and a slightly rising demand over the year (b). In contrast, the DFN measurements exhibit peaks caused by single demands.

at no cost.

The traffic values for all networks have been scaled to obtain comparable maximum total demand values (1 Tbps, 3 Tbps and 5 Tbps). To get undirected demands between nodes i and j , we considered the maximum of the two corresponding directed demands. In order to be consistent with the 15-minute traffic matrices of Géant, we aggregated the 5-minute traffic matrices of Abilene and Germany17 networks to 15-minute traffic matrices by taking the maximum value for each demand over the whole aggregation interval.

As shown in Fig. 4.2(a), the DFN matrices have a centralized structure with a large demand emanating from Frankfurt, which is a large entry point for cross-atlantic traffic, see [164]. They also exhibit temporal peaks caused by single academic institutions sending large amounts of traffic to another institution or to an international backbone. Therefore we also evaluated the energy savings in Germany17 with demand matrices generated using the Dwivedi-WaGner (DWG) model [174] based on population statistics. The resulting demands are much less centralized (compare Fig. 4.2(a) and (b), the area of each node is proportional to its emanating demand). The DWG model distinguishes between data, voice, and video traffic and computes demand values between two cities according to their distance and their number of inhabitants, employees, or households depending on the traffic class. From the single traffic matrix B_{SBN} obtained from the DWG model, we generated demand matrices for all time periods by applying the relative demand changes in the DFN measurements to the computed DWG matrix as follows. Given the DFN demands $D(t)$ over time, the maximum DFN demands D_{SBN} , and the static DWG demands b^{ab} , we calculate dynamic DWG demands $b^{ab}(t)$ in the following way:

$$b^{ab}(t) = b^{ab} \cdot d^{ab}(t) / d_{SBN}^{ab}.$$

The time-dependent scaling factor $d^{ab}(t) / d_{SBN}^{ab}$ takes values in the interval $[0, 1]$ normalizing the maximum DFN demand for every (a, b) . It hence rules out the domination effects caused by single demands in the measurements. Fig. 4.3 illustrates this effect for the daily total demand values over

Table 4.1: Network topologies and traffic [4].

Network	Nodes	Links	Time granularity	Reference
Abilene	12	15	every 15 min of 2005/05/08, every day of 2004/05, every two weeks between 2004/05/01 and 2004/07/24	[129]
Géant	22	36	every 15 min over 2005/05/05, every day over 2005/05/05–2005/06/04	[130]
Germany17	17	26	every 15 min of 2005/02/15, every day of 2005/02, every month of 2004 using DFN-measurements and Dwivedi-WaGner model	[131, 132, 174]

a month (a), and monthly total demand values over a year (b). The topologies and traffic data used in this chapter are summarized in Table 4.1.

4.3.2 Power and CapEx costs

In the following we briefly describe the network elements we used to design the IP-over-WDM architecture (see Fig. 2.1). Every network node can be equipped with one out of 13 different IP routers accommodating 16–208 line cards with a capacity of 640–8320 Gbps. Routers with a capacity of more than 640 Gbps are multi-chassis configurations incurring a multi-chassis setup cost such that the CapEx cost for the smallest router with 16 slots is $C^1 = 16.67$ and for the remaining 12 routers with 32 to 208 slots it roughly holds that $C^i = 111.67 + (i - 2) \cdot 29.17$ with $i \in \{2, \dots, 13\}$, also see [2]. Power consumption of routers can be calculated correspondingly assuming 2920 W consumed by a Line Card Shelf (LCS) and 9100 W consumed by a Fabric Card Shelf (FCS) (see Section D.1 in Appendix D). The corresponding CapEx and power costs of the routers used in this study are summarized in Table 4.2.

We considered a 40 Gbps colored line card interface that connects the IP router to the WDM system. To estimate the cost of this interface following [2], we combined a 40 Gbps IP router slot card, a 4x10 Gigabit Ethernet (GbE) port card, and a 4x10G Extended Long Haul (ELH) muxponder at the cost of 19.42. A lightpath is set up using these interfaces at both ends of the path, hence $C^{LC} = 38.84$. The power was evaluated by combining a Cisco 4-port 10-GE Tunable WDMPHY Physical Layer Interface Module (PLIM) and a Modular Services Card which together consume 500 W [5, 7] and hence 1000 W for a single (undirected) lightpath.

We assume an 80-channel optical system. Following [2], an optical fiber installed on a physical link is composed of Optical Line Amplifiers (OLAs), Dynamic Gain Equalizers (DGEs), Dispersion-Compensating Fibers (DCFs), and WDM multiplexers. As in [2] we assume an OXC to be composed of Wavelength Selective Switches (WSSs), which results in a fixed cost and a cost that linearly scales with number of connected fibers. We may hence map the latter to the cost of fibers. The corresponding total cost C^e of a fiber depends on the length of the actual physical link. In our case it holds that $C^e \in [21.16, 31.83]$, $C^e \in [24.67, 179.16]$, and $C^e \in [20.85, 134.63]$ for the networks Germany17, Géant, and Abilene, respectively.

Table 4.2: Set of available routers [2, 5].

SHELF TYPE	CAPACITY [Mbit/s]	CapEx [unit]	POWER [kW]
SH-IP-640	640000	16.67	2.92
SH-IP-1280	1280000	111.67	14.94
SH-IP-1920	1920000	140.83	17.86
SH-IP-2560	2560000	170	20.78
SH-IP-3200	3200000	199.17	23.7
SH-IP-3840	3840000	228.33	26.62
SH-IP-4480	4480000	257.5	29.54
SH-IP-5120	5120000	286.67	32.46
SH-IP-5760	5760000	315.83	35.38
SH-IP-6400	6400000	398.33	47.40
SH-IP-7040	7040000	427.5	50.32
SH-IP-7680	7680000	456.66	53.24
SH-IP-8320	8320000	485.83	56.16
SH-IP-8960	8960000	515	59.08
SH-IP-9600	9600000	544.17	62
SH-IP-10240	10240000	573.34	64.92
SH-IP-10880	10880000	602.51	67.84
SH-IP-11520	11520000	631.68	70.76
SH-IP-12160	12160000	714.18	82.78
SH-IP-12800	12800000	743.35	85.7
SH-IP-13440	13440000	772.52	88.62
SH-IP-14080	14080000	801.69	91.54
SH-IP-14720	14720000	830.86	94.46

Table 4.3: Cost of optical components and of line cards [2, 5, 6].

TYPE		DETAILS		CapEx [unit]	Power [W]
IP/MultiProtocol Switching (MPLS) Slot Card	Label Router	40 Gbps capacity, 1 slot/1 slot		9.17	350
IP/MPLS Card	Router Port	4 x 10 Gigabit Ethernet, Long Reach (1550 nm, 80 km reach), 1 slot occupied (40 Gbps)		4.20	150
WDM muxponder		10G x 4, ELH (1500 km)		6.05	colored interface without muxponder
Optical (OLA)	Line Amplifier	ELH ($R = 80$ km span)		2.77	not considered
Dynamic (DGE)	Gain Equalizer	80 channel systems		3.17	0
WDM (multiplexer/de-multiplexer/receiver)	Terminals + booster/amplifier	80 channel, (Long Haul (LH), ELH, Ultra Long Haul (ULH))		10.83	not considered
Dispersion-Compensating Fiber (DCF)		ELH (1500 km reach) with costs related to OXCs, L^e denotes the length (in km) of physical supply link $e \in E$		$2 \cdot 10.42 + 0.0091 \cdot L^e$	0

4.4 Results

This section describes computational results on the data presented in the previous section. We first explain the essence of our results using the Germany17 network and the 96 DFN traffic matrices given for every 15 minutes of 24 hours, and discuss these results in detail. Thereafter, we illustrate that for DUFL and DUDL these results are invariant against changes of

- the ratio between demand values (scaled to 1, 3 and 5 Tbps maximum total demand) and the capacity granularity (40 Gbps lightpath bitrate),
- the network (Germany17, Abilene, and Géant),
- the time scale of the demands (every 15 minutes over a day, every day over a month, and every month over a year), and

Table 4.4: Size of the MILPs to compute the SBN in terms of variables and constraints. Simple upper or lower bound constraints for variables are not counted.

Network	variables			constraints
	bin	\mathbb{Z}_+	\mathbb{R}_+	
Abilene	156	87	1452	237
Germany17	221	6559	4352	468
Géant	286	3006	9702	773

- the structure and spatial distribution of the demand matrix (DFN measurements or Dwivedi-Wagner model using population statistics).

while we report on the influence of the maximum total demand and the spatial distribution of traffic (measurements vs. population statistics) on the success of FUFL. All occurring MILPs have been solved using CPLEX 12.1 [175] as a black-box solver with a time-limit of 1 hour on a 64-bit Intel 3.00 GHz Central Processing Unit (CPU) with 8 GB main memory. The size of the MILPs to compute the SBN in terms of variables and constraints is presented in Table 4.4. These correspond to the size of model (4.3). Notice that the size of the models (B.1) and (B.2) cannot be explicitly stated since it very much depends on the concrete solution for the SBN. The number of variables however is always reduced since the variables for IP routers and fibers are fixed. For DUFL even most of the path-variables can be omitted if the paths are not used in the SBN solution. Also the number of constraints in (B.1) and (B.2) is bounded from above by the number of constraints in (4.3). We do not count simple bound constraints such as non-negativity constraints or (B.1e).

4.4.1 Results for the Germany17 network with DFN matrices over one day

We could compute an energy-minimal solution within seconds or minutes for most of the DUFL scenarios on the Germany17 network with the 15-minute DFN matrices for a day. For only a few instances we hit the time-limit with an optimality gap (relative difference between the number of line cards in the best solution and a mathematically proven lower bound to this number) below 5%. The optimization problem corresponding to DUDL is harder to solve. All DUDL runs hit the time-limit with optimality gaps of 11%–30% (1 Tbps), 6%–25% (3 Tbps), and 3%–15% (5 Tbps). A higher relative optimality gap for DUDL can be observed for the 1 Tbps maximum total demand than for 5 Tbps. All comparisons of the three strategies are made against the lower bound on the number of line cards in use (marked in the figures as DUDL-lower), which corresponds to an upper bound on the maximum possible energy savings in the considered scenario. Note that there are no dual bounds for FUFL since no optimization is performed with this approach. For almost all DUFL runs dual bounds and primal solution values are identical, which means that the computed solutions are optimal. In the following figures we hence removed the dual bound for DUFL since it cannot be distinguished from the (primal bound) solution curve. Whenever reporting on power values, we consider the total power consumptions of all active line cards assuming a value of 500 W for every single line card as explained in Section 4.3.2.




Figure 4.4: The figures show the power consumption in kilowatt with the three strategies on the Germany17 network, DFN traffic, 1/3/5 Tbps, every 15 minutes on February 15, 2005. The difference between FUFL and DUFL is much larger than the additional benefit of DUDL. Note that the y-axes are not identical.

Fig. 4.4 illustrates the (close to) optimal power consumptions obtained with FUFL and DUFL as well as the power consumptions and dual bounds for DUDL for each of the three demand scalings. All the proposed approaches make use of the dynamics of traffic and follow the total demand curve (compare with Fig. 2.7(b)). The energy savings achieved with DUFL and DUDL are nearly identical and much larger than with FUFL. The flexibility of DUFL to reroute traffic saves a significant amount of energy compared to FUFL. In contrast, reconfiguring the logical topology in the physical layer (DUDL) does not give much additional profit. In the 5 Tbps scenario the DUFL and DUDL curves nearly coincide. The lower bound for DUDL proves that only a small amount of energy can be saved compared to DUFL. There seems to be more tolerance in the 1 Tbps scenario. In this case we cannot verify whether our DUDL solutions are optimal or whether solutions closer to the lower bound exist. The constant periods of power consumption in Fig. 4.4(a) correspond to the minimal number of line cards that are needed to maintain IP connectivity.

More precisely, in the 1 Tbps scenario, the line cards of the network consume over the day 0.89/0.75/0.50/0.38 MWh for SBN/FUFL/DUFL/DUDL, respectively. The corresponding values for 3 Tbps and 5 Tbps are 2.11/1.52/0.94/0.82 MWh and 3.29/2.15/1.36/1.28 MWh. Notice that

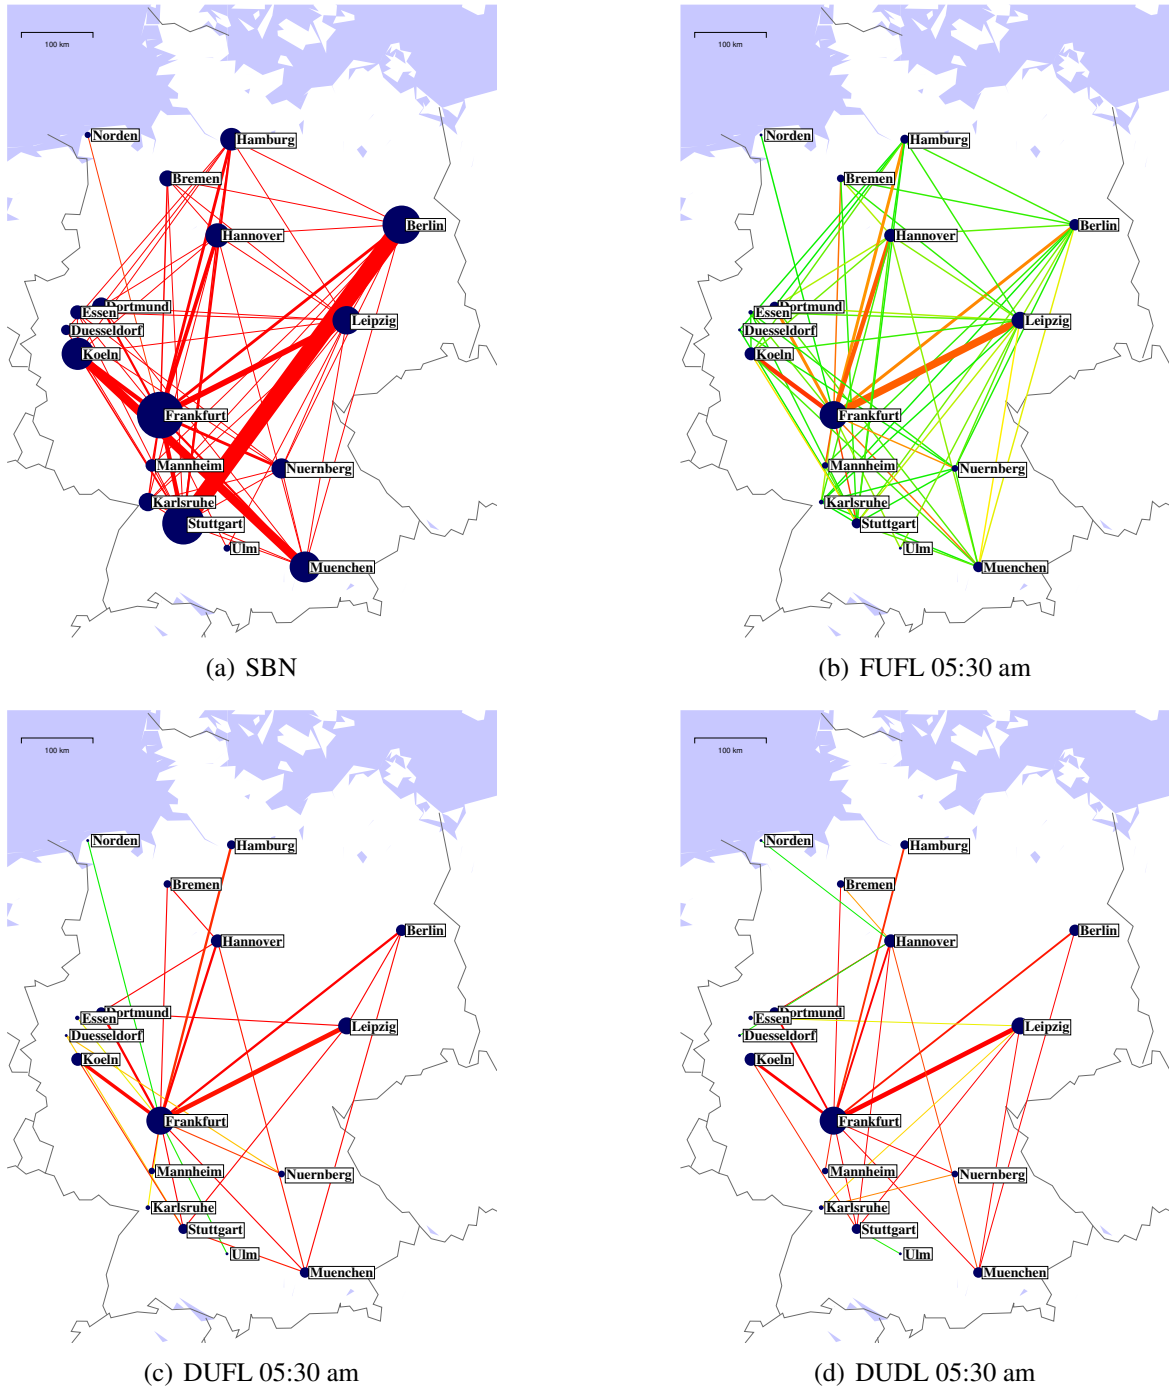


Figure 4.5: Logical topologies of the SBN and of the computed solutions with FUFL, DUFL, and DUDL at 05:30 am for the Germany17 network with DFN measurements, maximum total demand 5 Tbps. The color of a link corresponds to the load (high (red) and low (green)). The width of a link refers to its capacity. The size of a node represents its emanating demand.

these values correspond to the power consumption of a single day. Accumulating these values over a year (multiplying by 365) results in power consumption of 1201/785/496/467 MWh over the year in the 5 Tbps scenario.

Although we focus on the comparison of the FUFL, DUFL and DUDL approaches, it may be interesting from the practical perspective to evaluate the energy savings against the SBN. One should however not overestimate these savings since the SBN may be overprovisioned (peak demands typically do not occur simultaneously). In the 3 Tbps scenario FUFL reduces the power of the active line cards in low-demand hours by up to 38% at 05:30 am (72% for DUFL and 77% for DUDL). Even in a high-load scenario the savings are significant (17%, 39%, and 44% for FUFL, DUFL and DUDL at 02:45 pm, respectively). Considering the power consumption at 05:30 am and 02:45 pm for a maximum total demand of 3 Tbps, 25% of power for FUFL, 55% for DUFL, and 59% for DUDL can be saved in the early morning compared to the peak hour.

We observe that already the simple approach FUFL allows to save substantial amount of energy. The savings however depend on the ratio of the maximum total demand and the capacity of a single WDM channel W . If this ratio is low (traffic demands are low compared to a coarse lightpath bitrate) there is little potential to save energy with FUFL since a single lightpath might be sufficient to transport demands between pairs of nodes. Such single lightpaths cannot be switched off without the flexibility of IP rerouting – IP connectivity has to be maintained also in very low-demand hours. On the other hand if the demands are very large there are potentially many active lightpaths serving the same pair of nodes. The traffic on such a logical link is likely to drop below the 40 Gbps (80 Gbps, 120 Gbps, ...) threshold. It follows that with a constant bitrate of 40 Gbps for the line cards and increasing demands (from 1 Tbps to 5 Tbps) the relative outcome of FUFL increases which can be observed in Fig. 4.4(a) – (c). Compare also Fig. 4.7(a) with 4.7(c) and Fig. 4.7(b) with 4.7(d).

For a low-demand hour at 05:30 am in the 5 Tbps scenario, Fig. 4.5(a) – (d) show the logical topologies corresponding to the SBN and the considered approaches FUFL, DUFL, and DUDL, respectively. Although FUFL allows to reduce the link capacity (number of active line cards), the logical topology remains nearly the same because existing logical links have to be maintained even for small amount of traffic. In contrast, the logical topology changes significantly with DUFL and DUDL. Moreover, lightpaths in DUFL and DUDL are highly utilized, as opposed to FUFL, see Fig. 4.6. Remember that DUFL may only use logical links that exist in the SBN in contrast to the energy saving scheme DUDL which may set up new lightpaths (e.g., Hannover–Norden, compare Fig. 4.5(a) and (d)). Nevertheless the logical topologies of DUFL and DUDL are very close to each other with a similar number of active line cards.

To understand the relatively poor outcome of DUDL, one has to consider two extreme scenarios. If the demand in the network is very large, the logical topology in the SBN is close to a full mesh (see Fig. 4.5(a)). Since DUFL may use any logical link from the SBN, the DUFL solution is (close to) optimal and DUDL cannot benefit from choosing lightpaths not existing in the SBN. If, on the other hand, the demands are very small, the optimal logical topology of the SBN will be a tree. Both DUFL and DUDL will find a tree network. These trees might differ, but they use the same number of line cards. Again DUDL cannot benefit compared to DUFL. For the success of DUFL it is also crucial that we allow splitting of traffic demands in the logical domain, which lets DUFL fill up the established lightpaths to a high extent. This is illustrated in Fig. 4.6(a) and Fig. 4.6(b). Moreover, the lack of power-hungry network elements in the WDM layer [34] leads to the lack of potential to save energy by rerouting of lightpaths.




Figure 4.6: Average logical link utilization over all active logical links, Germany17, DFN, every 15 minutes. DUDL and DUFL achieve high lightpath utilization in contrast to FUFL.

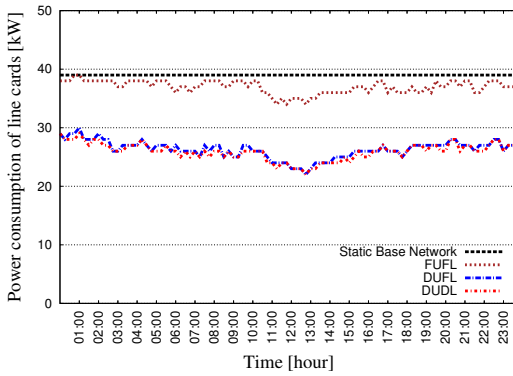
4.4.2 Varying input parameters

To make sure that the observed results do not depend on specific assumptions on the network or the input data we evaluated the performance of FUFL, DUFL, and DUDL for various parameter combinations of the network, the time scale, and the demand scalings and patterns. Since the observed results were basically consistent over all these scenarios, we show an example diagram for each of the variations. In the previous section, we have already shown that our results are nearly invariant against variations of the ratio between the demand and the capacity granularity, compare Fig. 4.4(a) – (c).


Varying the network Secondly, we varied the network topology by using the traffic measurements on the Géant and Abilene networks described in Section 2.5. Fig. 4.7 shows the power consumption curves over time for these two networks. The difference in the outcome of the three approaches is similar to the results observed for Germany17. The success of FUFL depends on the size of the demands while DUFL performs constantly well with almost no additional benefit by using DUDL. Notice that with the Abilene network, all obtained network topologies and power consumption values are optimal.

Varying the time scale Thirdly, we varied the time scale, which changed the structure of the demand variations over time. For the DFN measurements on the Germany17 network, we not only considered the 15-minute demand matrices over a day, but also aggregated measurements for every day over a month, and for every month over a year, as explained in detail in Section 2.5. Fig. 4.8 shows the power consumption over time with FUFL, DUFL, and DUDL for the latter two time scales on the Germany17 network with DFN demands. One can see that DUFL and DUDL are very close to each other and can save much more energy than FUFL also on these time scales.


Varying the demand pattern Eventually, we varied the structure of the demands by using demands generated with the Dwivedi-WaGner (DWG) model instead of the DFN measurements. Fig. 4.9 shows the power consumption over time with the DWG demands on the Germany17 network. It turns out that the essential result does not change. The benefit of DUFL compared to FUFL is very large, and DUFL and DUDL differ only marginally. The plotted dual bounds show that the network topologies obtained in each time slot are close to the optimum. In particular, this




(a) Abilene network, 1 Tbps



(b) Géant network, 1 Tbps



(c) Abilene network, 5 Tbps



(d) Géant network, 5 Tbps

Figure 4.7: The results are independent of the network topology. The figures show the power consumption over time with the three approaches on the Géant and Abilene network with measured traffic demands, scaled to 1 Tbps and 5 Tbps of maximum total demand, every 15 minutes of a day.

shows us that this effect is due to structural differences of the considered rerouting concepts and not the result of some heuristic solution procedure.

On the other hand, it can be seen that the power reduction compared to the SBN using the FUF approach is marginal, which is in contrast to the results for the DFN-demands (compare Fig. 4.4(b) with Fig. 4.9). Recall that the DFN-traffic is very centralized with a large concentration of demand in Frankfurt, while the DWG-demands are more evenly distributed (compare Fig. 4.2(a) with Fig. 4.2(b)). This has mainly two consequences. First, the line cards are more evenly distributed among the nodes in the 3 Tbps DWG solution with a minimum of 2 line cards (Norden) and a maximum number of 15 line cards (Frankfurt, Leipzig, and Muenchen) compared to a minimum of 1 (Norden and Ulm) and a maximum of 32 (Frankfurt) in the 3 Tbps DFN SBN solution. And secondly, the number of used physical paths is larger in the DWG solution compared to the DFN solution, which influences the number of parallel lightpaths. In the DWG SBN 87 physical paths are in use with a maximum number of 2 lightpaths using the same physical path. The DFN solution uses a total of 50 paths with a maximum of 14 channels on the same path. Note however that the total number of lightpaths in use is almost the same in both scenarios (88 for DFN and

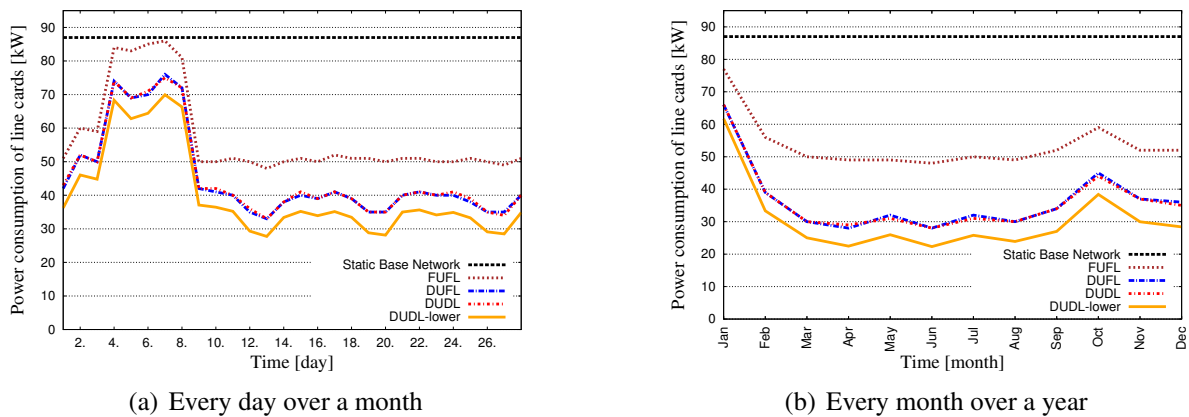


Figure 4.8: The results are independent of the time scale. The figures show the power consumption over time with FUFL, DUFL, and DUDL in the Germany17 network, DFN traffic, 3 Tbps, for a month and a year. The meaning of the curves is the same as in the 15-minute Fig. 4.4(b).




Figure 4.9: Power consumption over time with FUFL, DUFL, and DUDL in the Germany17 network, 3 Tbps, DWG traffic, 15 minutes over one day. The change of the spatial distribution of traffic affects the success of FUFL. Compare with Fig. 4.4(b).

89 for DWG), because the sum of all demands is identical (3 Tbps). Since FUFL, as explained in Section 4.2, may switch off line cards only in the presence of parallel channels, the power consumption can only slightly be reduced with DWG demands. Here the spatial distribution of the traffic influences the impact of FUFL approach on energy saving. In contrast, we do not observe any differences for DUFL and DUDL.

4.5 Conclusion

Significant energy savings can be potentially achieved with all the considered approaches FUFL, DUFL and DUDL. DUFL contributes the most to the energy savings. The additional flexibility of routing in the WDM layer in DUDL does not bring much further savings compared to DUFL. This is caused by densely interconnected logical topology of the SBN, the assumption of multipath

routing, and the spatial distribution of traffic demands which does not change much over time.

Although FUFL does not bring as much benefits as DUFL, it is much simpler and can be run locally. The decisions to switch line cards on and off are made without the global knowledge of the network and its load. We observed however that the success of the simple approach FUFL depends on the capacity of a WDM channel compared to the size of the demands as well as on the regional distribution of the demand. Therefore it is of interest to network operators [18]. Given a constant lightpath capacity, the benefit of FUFL increases with increasing demands. Given the same lightpath capacity, the same total traffic and the same dynamics over time, less energy can be saved with FUFL under (non-realistic) evenly distributed demands over space (DWG) than under demands with distribution originating from measurements (DFN).

5 A step towards bringing energy saving with EA-ARSs into reality

Results from Chapter 4 show that there is a significant potential of energy saving that can be achieved with switching line cards off in low-demand hours. However, how much energy saving out of this potential can be achieved in reality? We list the constraints that the Energy-Aware Adaptive Routing Solutions (EA-ARSs) have to face in reality. Moreover, a heuristic is proposed as an alternative to the Mixed-Integer Linear Programming (MILP) model from Chapter 4, and evaluated on a set of realistic scenarios with comparison to two alternative heuristics. Parts of this chapter have been published in [112, 137, 176].

5.1 Evaluation criteria for EA-ARSs

A network using an EA-ARS needs to deal with several constraints during its operation. We clearly identify them in this section as a roadmap towards evaluation of the energy saving mechanisms proposed in the literature, and as indication for the standardization bodies towards bringing energy saving into reality of backbone networks. Please mind that the constraints partly overlap with the taxonomy presented in Section 3.2.1. Apart from the classical physical layer constraints, and the constraints on installed devices, the following issues need to be considered by the algorithm dynamically reducing power of Internet Protocol (IP)-over-Wavelength Division Multiplexing (WDM) networks.

1. **Targeted devices:** First of all, the devices that are targeted to be switched off (or put into standby mode) have to support such functionality [3]. It must be possible to remotely change the state of the devices. Moreover, the difference between power consumption of a device in an on state and power consumption of a device in a standby state must be high enough so that the dynamic toggling of the states pays off. Eventually, the devices must be tested against the failure rates, since frequent toggling of the device state is expected to increase its Mean Time Between Failures (MTBF) [177].
2. **Physical layer constraints:** Classical constraints cannot be ignored by the energy saving mechanisms. In particular, maximum length of a lightpath cannot be exceeded. Availability of a wavelength on the entire route of the lightpath needs to be validated before its establishment (Routing and Wavelength Assignment (RWA)), or wavelength conversion must be ensured.
3. **Constraints on installed devices:** In contrast to network design, devices cannot be dynamically installed and deinstalled during network operation. Therefore the energy saving algorithms must make sure that the devices required in the provided solution are present and

available in the network. As an example, an EA-ARS cannot require 17 active line cards in a node which is equipped only with 16 line cards.

4. **Impact on Quality of Service (QoS):** Switching off devices decreases capacity of the network leading to lower power consumption. However, the dynamic reconfiguration comes at the cost of establishing/releasing lightpaths, potential rerouting or IP traffic demands, longer paths that the traffic demands take from source to target nodes, and eventually the risk of congestion once traffic unexpectedly increases to a high extent. All these costs can be mapped into the general term QoS. Therefore the EA-ARSs should be aware of the fact that energy saving influences the QoS, and try to find a trade-off between the two. Overprovisioning of the network cannot be avoided, however it is expected to have different behavior than today, because it will be dynamic, i.e., the level of overprovisioning will be adequate to the current traffic demands.
5. **Computation time:** Traffic dynamically changes over time. While the traffic burstiness is much lower in the backbone networks than in the access networks, the reaction of a EA-ARS to the changing traffic conditions should be limited in time. Therefore the runtime of an EA-ARS should be sufficiently small in order to be able to follow the changes of traffic.
6. **Triggering events:** The aim of the EA-ARSs is to dynamically adapt the network configuration to the changing traffic conditions. However, the reconfiguration should on one hand take place frequently enough to follow the traffic conditions and save sufficient amount of energy, and on the other hand not too often in order not to introduce instabilities into the network. Therefore clear triggering events (such as violation of load levels on certain links, or changing level of total load) need to be defined. Periodic execution of the EA-ARS may result in a waste of computation resources when traffic does not change, or overlooking of the sudden traffic change.
7. **Operation and network knowledge:** Each EA-ARS operates executed either at each network node, or centrally. The difference is mainly about whether the input data for the EA-ARS needs to be delivered to one node (centralized operation), or to all nodes (distributed operation). Moreover, the solution needs to be delivered back to the nodes in the centralized operation.

No matter whether the operation is centralized or distributed, the required input data (referred to as network knowledge) may be either global or local. The following aspects constitute the network knowledge:

- a) Complete traffic matrix
- b) Routing of IP traffic demands over the logical topology
- c) Physical topology and installed devices
- d) Available wavelengths and routing of lightpaths (including physical layer constraints)
- e) Load on the logical link (which can be computed from 1 and 2)

Complete or partial network knowledge may be required by an EA-ARS.

8. **Protection consideration:** Reliable backbone is an important issue for network operators since backbone links have high capacities. Energy saving has to be therefore traded against protection level, similarly to QoS. Therefore, EA-ARs should take into account the level of protection when switching off devices for energy saving.
9. **Reconfiguration costs:** Several optimal network configurations may exist with respect to power consumption. However, the transition from an optimal configuration in one time period to an optimal configuration in another period may have different costs in terms of removed lightpaths, their rerouting, and IP rerouting (including different splitting of single flows when multipath routing is allowed). In an extreme case, two feasible optimal network configurations may exist for the two subsequent time periods, however the transition from one configuration to another configuration can be infeasible without traffic loss. Therefore the limitation of the differences between network configurations for the subsequent time periods is also an important challenge for the EA-ARs.
10. **Future traffic assumption:** It is obvious that future traffic is unknown. Even if the computation is triggered using up-to-date traffic measurements, traffic burstiness in the core is limited and future traffic increase or decrease is estimated based on historical data, traditional overprovisioning is still needed to cover the unexpected traffic increase. The assumption of known future traffic by the EA-ARs corresponds to an upper bound of energy savings.

Differently from Chapter 4 and most of the related work, we consider a set of time periods T consisting of past and future time periods (T_{past} and T_{fut}) in this chapter. T_{past} is used for the design of Static Base Network (SBN) determining the devices installed in the network and their configuration (compare with Section 4.1). T_{fut} is used for evaluation of energy saving during network operation with an EA-ARS. Therefore $T_{past} \cap T_{fut} = \emptyset$, and $T = T_{past} \cup T_{fut}$. It addresses the issue that traffic that will flow through the network is unknown during its design stage, what is the case in reality.

5.2 Energy Watermark Algorithm

The proposed heuristic for dynamic network reconfiguration to save energy in an IP-over-WDM network focuses on the IP layer, and uses the following variables defined over $t \in T_{fut}$. The flow variables $f_{ij}^{ab}(t) \in \{0, 1\}$ denote whether the traffic demand originated at node a and targeted to node b traverses the logical link from i to j at time t . Differently to the MILP formulations from Chapter 4, Energy Watermark Algorithm (EWA) assumes single-, shortest-path routing of traffic demands over the logical topology. Similarly to the $y_p(t)$ in the MILPs from Chapter 4, the variables $y_l(t) \in \mathbb{Z}_+$ determine the number of lightpaths used to realize the logical link $l \in L$ at time t , which in turn determine the powered on line cards. Finally, $x_i^{LC}(t) \in \mathbb{Z}_+$ is the number of line cards powered on at each node i at time t , which is bounded by the number of installed line cards X_i^{LC} in each node of the SBN.

Power consumed by a line card, Line Card Shelf (LCS) and Fabric Card Shelf (FCS) is denoted by \mathcal{C}^{LC} , \mathcal{C}^{LCS} and \mathcal{C}^{FCS} , respectively.

We define the network configuration as the set of network nodes V with installed line cards X_i^{LC} (powered on or off), established lightpaths forming logical links $y_l(t)$ and IP routing of traffic

demands $f_{ij}^{ab}(t)$. Differently from Chapter 4, we use directed graphs and traffic demands in this chapter.

5.2.1 General idea

The EWA adapts the network to current traffic demands in order to save energy on one hand, and limit the load on logical links in order to ensure certain QoS on the other hand. EWA uses a low and a high watermark (W_L and W_H) defined as thresholds on the utilization of the last lightpath on a logical link. Exceeding the W_H triggers attempts to establish additional lightpath(s) in order to avoid overload of the network. Exceeding the W_L triggers attempts to release lightpath(s) in order to switch off idle line cards and save energy. EWA makes sure that the maximum utilization of last lightpath on a logical link ψ is not exceeded when trying to release lightpaths.

Alg. 1 shows the main pseudocode of EWA. Details of its subroutines are presented in the next section. The algorithm takes as input the network configuration in previous time period $t - 1$ (network nodes V with installed line cards X_i^{LC} , established lightpaths forming logical links $y_l(t - 1)$ and IP routing of traffic demands $f_{ij}^{ab}(t - 1)$), traffic matrix $D(t)$ for the current time period t , capacity of a lightpath (WDM channel) W , W_L , W_H and ψ . Updated network configuration is returned as output of the algorithm.

EWA first checks whether all the demands in the current network configuration are routable, and iteratively tries to establish additional lightpath(s) for the unroutable demands (if any), starting from the largest ones (line 1). The logical links on which watermarks are exceeded are identified next (line 2), and violation of the W_H is checked, starting from the logical links with the highest utilization of the last lightpath (line 3). For each overloaded logical link (from the most overloaded to the least overloaded), the algorithm first tries to increase the capacity of the logical link if a demand with the same source and target flows through it. If this is not the case, attempts are made to establish lightpath(s) for the possibly biggest demand flowing through the overloaded logical link.

Once load is lower than W_H for all logical links, or it is impossible to reduce overload anymore, violation of the W_L is checked starting from the least loaded logical links (line 4). One lightpath per iteration is tried to be released making sure that ψ is not exceeded.

5.2.2 Main steps

We then summarize the main steps of EWA. Details of these steps are presented in Appendix C.

Algorithm 1 Pseudo-code of EWA.

Input: netConfig from period $t - 1$, current Traffic Matrix (TM) $D(t)$, W , W_L , W_H , ψ

Output: Updated netConfig

- 1: ensureDemandsRoutability(netConfig, $D(t)$, W);
 - 2: sortedLLsExceedingWMs = getSortedLLsExceedingWMs(netConfig, $D(t)$, W_L , W_H);
 - 3: establishNecessaryLpaths(netConfig, $D(t)$, W , sortedLLsExceedingWMs, W_L , W_H);
 - 4: releaseUnnecessaryLpaths(netConfig, $D(t)$, sortedLLsExceedingWMs, W_L , W_H , ψ);
-

5.2.2.1 Ensuring Routability of Demands

The main requirement on an energy-saving algorithm is not to influence the connectivity of the network. This means that given a set of traffic demands $D(t)$, there must be at least one path in the logical topology determined by $y_l(t)$ to route each traffic demand. Therefore the first step of EWA is to ensure that all the demands are routable. In particular, starting from the current network configuration the unroutable demands are identified and sorted in descending order according to the demand value. Then, iteratively for each unroutable demand, a logical link is tried to be established. The routine iterates until all the demands are routable or it is impossible to establish additional logical links. Finally, updated network configuration is returned.

5.2.2.2 Establishing Necessary Lightpaths

During this step the algorithm ensures that the capacity of the network is able to satisfy the current traffic demands $D(t)$. In particular, given the updated network configuration and the sorted list of logical links where the high watermark is exceeded, the algorithm iteratively tries to add lightpaths. Initially, the set of logical links is ordered in descending order of utilization of the last lightpath. Then, for each logical link violating the high watermark, attempts to establish new lightpath(s) are made. The algorithm tries to offload the current logical link by establishing additional lightpaths and shifting the demands from the current logical link to the new lightpaths. In particular, two cases are identified based on the demands flowing through the current logical link: a) if a demand with the same source and target nodes as the current logical link flows through the current logical link, and its value is higher than the capacity of the logical link, parallel lightpaths are added to the logical link (subject to availability of transmitters and receivers), until the current demand is satisfied, b) otherwise, a new logical link is established directly between the source and destination of a demand flowing through the current logical link if transmitter(s) and receiver(s) are available. Note that the demands satisfying condition b) are sorted in decreasing order, starting from the biggest one. The idea is to fill the newly established lightpaths to a high extent (preserving the W_H), which corresponds to efficient usage of the power of the corresponding line cards. The algorithm iterates over the set of demands until the high watermark constraint is ensured (or no more demands remain). The procedure is repeated for all logical links violating the high watermark.

5.2.2.3 Releasing Unnecessary Lightpaths

The last step of the EWA is to attempt to release unnecessary lightpaths in order to save energy. This subroutine takes as input the current network configuration and the set of logical links violating the low watermark. Then, the algorithm iteratively selects a logical link, starting from the one having the lowest utilization of the last lightpath. A single lightpath of the current logical link is attempted to be released keeping the constraints on routability as well as the constraints on ψ in the whole network. Then, the procedure is repeated for all logical link violating W_L , until it is possible to release lightpaths.

5.2.3 Differences to related work

To the best of our knowledge, no heuristic based on watermarks triggering establishing and releasing of lightpaths has been investigated so far in the literature focusing on power consumption.

EWA was inspired by [178]. The main difference to [178] (apart from the focus) is that EWA does not forbid releasing lightpath(s) in the case when some lightpath(s) have already been established in the same run of the algorithm. This allows more attempts to release lightpaths (more aggressive energy saving) and is not critical for the operation of the network, since releasing of lightpaths is performed in the last step of the algorithm (line 4 of Alg. 1). Lightpaths are released only after establishing the new ones, so that the necessary rerouting can be performed in a controlled way. Moreover, differently from [178], EWA is able to add or delete more than one lightpath during the execution of the algorithm. Preliminary investigations have shown that adding or deleting of a single lightpath is insufficient to adapt the network to the traffic changes based on input data originating from measurements (Section 2.5.2 and [4]). Finally, the parameter ψ is added in order to trade between QoS and energy saving.

Another work which uses the concept of watermarks is [179]. The authors of [179] introduce the load balance indicator (bound on maximum lightpath load of the new logical topology at reconfiguration point), which is similar to ψ , however defined on a lightpath, and not on (the last lightpath of) a logical link. Moreover, the watermarks are used only to trigger solving the optimization problem, which does not use the watermarks themselves, but only the load balance indicator. The optimization problem gives the possibility to re-distribute the load in the network by rerouting only, without adding or deleting lightpaths, however solving an optimization problem is time consuming and therefore impractical in real operation. The authors of [179] propose also a so-called softbound approach, where reconfiguration is done only after three observation periods since the previous reconfiguration if in between there is no violation of the high watermark. Eventually an approximate mathematical model is proposed in [179], where the added and deleted lightpath(s) are chosen among some candidate lightpaths. Algorithms for the selection of the candidate lightpaths are also described in [179]. They perform the routing and wavelength assignment in order to reduce the complexity of the mathematical formulation.

5.3 Alternative heuristics

5.3.1 Least Flow Algorithm

The Least Flow Algorithm (LFA) was proposed by Luca Chiaraviglio et al., and targets the minimization of the number of logical links powered on to satisfy a given traffic demand, adopting a modified version of the Least-Flow algorithm from [61], which is designed to work in IP networks. In particular, the logical links with the lowest amount of traffic flowing on them are targeted first. The intuition is that it is simpler to switch off links carrying a low amount of flow rather than a link whose flow is close to the link capacity. Alg. 2 reports a schematic description of the LFA heuristic. In particular, the set Y_p^{SBN} of logical links from the SBN, the current traffic matrix $D(t)$ and the maximum link utilization δ are provided as input. At the beginning, the logical links are sorted with increasing flow (line 1). Then, at each iteration, the considered link is removed from the topology (line 3), and traffic is then rerouted on the residual topology (line 4-5). After rerouting, if connectivity and utilization constraints are still fulfilled (line 6), then the selected link is definitively powered off. Otherwise it is left on. The procedure is repeated for all the logical links (line 2). In the case that the current traffic matrix cannot be satisfied by the SBN, LFA does not power off any link.

Algorithm 2 Pseudo-code of LFA.

Input: Set Y_p^{SBN} from SBN, current TM $D(t)$, maximum utilization δ

Output: Updated netConfig

```

1: LLS=sortLeastFlow( $Y_p^{SBN}$ );
2: for  $j = 1; j \leq \text{size}(\text{LLS}); j++$  do
3:   disableLogicalLink(LLS[j]);
4:   paths=computeAllShortestPaths( $D(t)$ );
5:   computeAllLinkFlow(paths,  $D(t)$ );
6:   if (checkPaths(paths) == false) ||
      (checkFlows(paths,  $\delta$ ) == false) then
7:     enableLogicalLink(LLS[j]);
8:   end if
9: end for

```

5.3.2 Genetic Algorithm

The Genetic Algorithm (GA) was introduced in [38] and implemented by Edoardo Bonetto et al., where it optimizes just the power consumption of the network. It has been adapted to directly minimize the power consumption and the reconfiguration cost. The GA is a meta-heuristic based on the principles of the natural evolution: the initial population evolves through several generations in which only a subset of the individuals survive from one generation to the following one. A set of individuals forms a population. A new individual (offspring) is created by modification and merging of two other individuals (parents).

Each individual represents a logical topology. Only feasible individuals can be part of the population. In particular, an individual is feasible if its corresponding logical topology satisfies the traffic matrix and the constraints on the maximum number of powered on devices. If the GA cannot find any feasible solution, the constraints are relaxed and dummy line cards are installed. However, the traffic routed over these additional line cards is considered as overload.

The individuals survive through generations depending on their fitness value which is defined as the weighted sum of normalized power consumption and normalized reconfiguration cost:

$$\min \left((1 - \alpha) \sum_{i \in V} x_i^{LC}(t) + \alpha \sum_{i \in V} \sum_{j \in V} \sum_{a \in V} \sum_{b \in V} r_{ij}^{ab}(t) \right) \quad (5.1)$$

where $\alpha \in [0, 1]$ denotes the cost associated with each unit of reconfigured traffic. The setting of α allows trading between power minimization (low α) and low reconfigured traffic (high α).

Alg. 3 describes briefly how the GA works. The algorithm requires as input the network configuration from previous time period $t - 1$ and the current traffic matrix $D(t)$. Moreover, three algorithm parameters have to be provided: the maximum number of generations without improvements M , the population size S and the offspring size κ . At the beginning, the first population is randomly generated (line 1). After this step the fitness is evaluated (line 3) and the evolution process begins (line 4). At each generation, the reproduction and the selection phases are repeated. In particular, new individuals are created during the reproduction phase (line 5), starting from two other individuals that have been randomly selected as parents. The reproduction phase includes

Algorithm 3 Pseudo-code of GA.

Input: netConfig from period $t - 1$, current TM $D(t)$, α , M , S , κ

Output: Updated netConfig

```

1: population=generateFirstPopulation( $D(t)$ , netConfig,  $S$ ,  $\alpha$ );
2:  $i=0$ ;
3: fitness=evaluateFitness(population,  $\alpha$ );
4: while ( $i \leq M$ ) do
5:   offspring=generateOffspring(population,  $D(t)$ , netConfig,  $\alpha$ ,  $\kappa$ );
6:   population=selectPopulation(population, offspring,  $\alpha$ ,  $S$ );
7:   newFitness=evaluateFitness(population,  $\alpha$ );
8:   if (newFitness  $\geq$  fitness) then
9:      $i++$ ;
10:  else
11:     $i=0$ ;
12:    fitness=newFitness;
13:  end if
14: end while
15: netConfig=applyConfiguration(population);

```

crossover and mutation concerning (sub)sets of logical links, that is merging parts of the parents and changing some logical links of the newly created individual (offspring). Then, a new population is created (line 6) by selecting the individuals with best fitness value from the old population and its offspring. At each generation, the best fitness value among all the individuals is selected and stored (line 7). If the fitness has not improved with respect to the previous generation, the current number of generations without improvement is incremented (line 9). The evolution process stops when the fitness value has not improved for more than a maximum number of generations (line 4). At the end, the individual with the best fitness is chosen, and the set of powered on lightpaths is updated (line 15).

5.4 Metrics

A trade-off between reconfigured traffic and the power consumption is investigated in this work. Therefore we look at the following metrics.

Power consumption (in general cost) of all line cards active in the network (as a function of time) is defined as:

$$P^{LC}(t) = C^{LC} \sum_{i \in V} x_i^{LC}(t) \quad (5.2)$$

The energy consumed by line cards over T_{fut} is defined then as:

$$E^{LC} = \sum_{t \in T_{fut}} P^{LC}(t) \Delta t \quad (5.3)$$

Cost (here power consumption) of active LCSs and FCSs is determined by the number of active

line cards. The number of LCSs used at each node is expressed as:

$$x_i^{LCS}(t) = \lceil x_i^{LC}(t)/W_{LCS} \rceil \quad (5.4)$$

where W_{LCS} is the capacity (in terms of line cards) of a LCS. The number of FCS is determined by the number of LCSs:

$$x_i^{FCS}(t) = \begin{cases} 0 & \text{if } x_i^{LCS}(t) \leq 1 \\ \lceil x_i^{LCS}(t)/W_{FCS} \rceil & \text{otherwise} \end{cases} \quad (5.5)$$

where W_{FCS} denotes the maximum number of LCSs that a FCS can interconnect. Cost (power consumption) of active LCSs in the network is given by:

$$P^{LCS}(t) = c^{LCS} \sum_{i \in V} x_i^{LCS}(t) \quad (5.6)$$

Similarly, the cost (power consumption) of active FCSs is defined as:

$$P^{FCS}(t) = c^{FCS} \sum_{i \in V} x_i^{FCS}(t) \quad (5.7)$$

The total cost (power consumption) of the whole network is hence defined as:

$$P^{TOT}(t) = P^{LC}(t) + P^{LCS}(t) + P^{FCS}(t) \quad (5.8)$$

and the total energy consumption of the whole network over T_{fut} is defined as:

$$E^{TOT} = \sum_{t \in T_{fut}} P^{TOT}(t) \Delta t \quad (5.9)$$

We consider the traffic that needs to be rerouted in order to reduce power consumption of the network. Therefore, let us define as $r_{ij}^{ab}(t) \in \mathbb{R}_+$ the amount of reconfigured traffic between a and b on the logical link from i to j at time t with respect to time $t - 1$, with $t > 0$ belonging to the set of future time periods T_{fut}

$$r_{ij}^{ab}(t) = \begin{cases} d^{ab}(t) \cdot f_{ij}^{ab}(t) - d^{ab}(t-1) \cdot f_{ij}^{ab}(t-1) & f_{ij}^{ab}(t) > f_{ij}^{ab}(t-1) \\ 0 & \text{otherwise} \end{cases} \quad (5.10)$$

We introduce the reconfiguration ratio over all subsequent pairs of time periods in T_{fut} with $t > 0$ as:

$$\xi = \frac{\sum_{t \in T_{fut}} \sum_{i \in V} \sum_{j \in V} \sum_{a \in V} \sum_{b \in V} r_{ij}^{ab}(t)}{\sum_{t \in T_{fut}} \sum_{a \in V} \sum_{b \in V} d^{ab}(t)} \quad (5.11)$$

This metric captures the amount of traffic which is reconfigured over all time periods in T_{fut} , normalized by the total amount of traffic which is exchanged in the network. The ξ may be greater than 1 since the reconfigured traffic is counted multiple times if it passes through multiple logical links from the source to the target.

Finally, we define in Eq. (5.12) the overload ratio ϕ to capture overload traffic in all periods $t \in T_{fut}$. ϕ should be as low as possible (ideally 0) to prevent service disruptions and loss of QoS. Note that we take into account each hop that a demand passes through when calculating the overload, even though traffic can be dropped only once in the network. Therefore, ϕ may also be greater than 1.

$$\phi = \frac{\sum_{t \in T_{fut}} \sum_{i \in V} \sum_{j \in V} \max \left(\sum_{a \in V, b \in V} d^{ab}(t) \cdot f_{ij}^{ab}(t) - \sum_{l \in L(i,j)} W_{yl}(t), 0 \right)}{\sum_{t \in T_{fut}} \sum_{a \in V} \sum_{b \in V} d^{ab}(t)} \quad (5.12)$$

Table 5.1: Considered networks and traffic data [4].

Network	Nodes	Phy. supply links (bidir.)	TMs considered at SBN design (T_{past})	TMs to evaluate energy savings (T_{fut})
Abilene	12	15	max. between 2004/07/01 and 2004/07/31 (granularity of 5 min.)	every 15 min. over 2004/08/27 (WD), 2004/08/28 (WE), 2004/08/29 (WE), 2004/09/02 (WD)
Géant	22	36	max. between 2005/05/05 and 2005/06/04 (granularity of 15 min.)	every 15 min. over 2005/06/07 (WD), 2005/06/10 (WD), 2005/06/11 (WE), 2005/06/12 (WE)
Germany17	17	26	max. between 2004/01 and 2004/12 (granularity of 1 month)	every 15 min. over 2005/02/15 (WD)

5.5 Evaluation scenario

We describe the network and traffic data used in this chapter. The same assumptions as in Section 4.3.2 on capacity of a lightpath W , power and cost values are used. The three physical supply network topologies Abilene, Géant and Germany17 are considered (see Figures 2.5(a), 2.5(b) and 2.7(a))

Traffic matrices with granularity Δt equal to 15 minutes are available for Géant, while for Abilene and Germany17 traffic matrices are measured every 5 minutes. In order to be consistent, we compute traffic matrices with granularity 15 minutes for Abilene and Germany17 by taking the maximum values out of the 3 corresponding 5-minute traffic matrices.

Total period covered by T_{past} (design of the SBN) equals 1 month for Abilene and Géant. Note that fine granular traffic matrices (5 min.) cover only one day in the Germany17 network [4]. Therefore we use 12 traffic matrices with time granularity equal to 1 month over the year 2004 to dimension the SBN. Since these traffic matrices contain average values of traffic demands over a month, they alter temporal peaks of the traffic, and hence higher overprovisioning is needed during the design of the SBN.

We consider the total period covered by T_{fut} equal to 1 day. We select the evaluation days belonging to two different types (Working Day (WD) and WeekEnd day (WE)) in order to compare the three EA-ARSSs under various traffic conditions.

The original traffic matrices from [4] are scaled in order to mimic actual traffic demands and to have comparable load in the three networks. We introduce the total demand per node $d_{|V|} = \sum_{(a,b) \in V \times V} d_{SBN}^{ab} / |V|$, and the corresponding unit Gigabit per second per node (Gpn). We consider three different load assumptions, and scale D_{SBN} to have the total demand per node equal to 100, 300 and 500 Gpn. The same scaling factors as for D_{SBN} are used for $D(t)$ for each $t \in T_{fut}$.

The networks, as well as the traffic data that we used for both the design of the SBN and evaluation of energy savings are summarized in Table 5.1. Total demand (defined as $\sum_{(a,b) \in V \times V} d^{ab}(t)$) over time is presented in Figures 5.1(a) and 5.1(b) for two exemplary days (2004/08/27 for Abilene and 2005/06/10 for Géant) and total load per node equal to 300 Gpn.

We use the GA [38] with the objective of Capital Expenditures (CapEx) minimization for the SBN design. The SBN is dimensioned to satisfy the maximum traffic matrix D_{SBN} , based on the

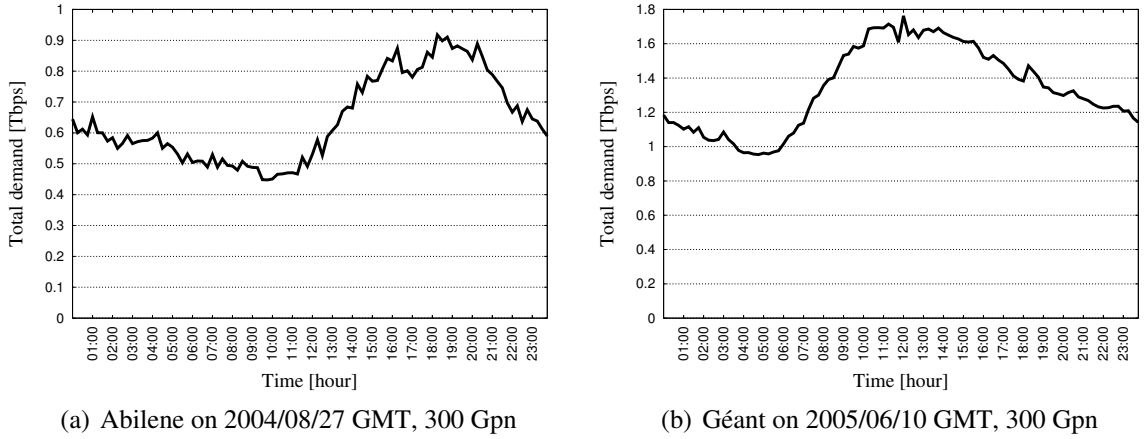


Figure 5.1: Total demand over time for the traffic data used for EWA evaluation of energy savings (total load per node 300 Gpn).

set of past time periods T_{past} :

$$d_{SBN}^{ab} = \max_{t \in T_{past}} d^{ab}(t), \quad \forall a, b \in V \quad (5.13)$$

This configuration may not be able to satisfy all future traffic demands, causing link overload and consequently dropping of traffic. For this reason, during the SBN design, we overprovision the network in such a way that the installed resources (line cards, LCSs and FCSs) may eventually sustain the future traffic. The overprovisioning factor γ is defined as the ratio between the light-path's capacity used during the SBN design and its full capacity. Thus, $\gamma \in (0, 1]$. For instance, with an overprovisioning factor of 0.4, the SBN is designed considering that lightpaths can use just 40% of their capacity, meaning that 60% more resources are installed with respect to the expected requirements.

The LFA, GA and EWA are executed at each t out of the set of future time periods $T_{fut} \subset T$.

5.6 Results

We have implemented the EWA algorithm on a custom simulator written in Java. The EWA simulations are run on a personal computer with a Dual Core Central Processing Unit (CPU) at 2.4 GHz and 2 GB of Random Access Memory (RAM). The alternative heuristics LFA and GA were implemented by Edoardo Bonetto and Luca Chiaraviglio on custom simulators written in C++. The LFA and GA simulations are run on a server machine equipped with 2 Quad CPU at 2.66 GHz and 4 GB of RAM. We focus on SBNs designed for $d_{|V|} = 300$ Gpn of traffic and overprovisioning $\gamma = 0.5$. Unless specified otherwise, we adopt the following set of algorithm parameters: $\delta = 1.0$ for LFA, $\alpha = 0.1$, $M = 500$, $S = 30$ and $\kappa = 20$ for GA, $W_L = 0.1$, $W_H = 0.9$ and $\psi = W_H$ for EWA. We first compare the algorithms with the following metrics: line cards daily energy consumption E^{LC} , total daily energy consumption E^{TOT} , reconfiguration ratio ξ , and overload ratio ϕ . Table 5.2 reports E^{LC} for the different scenarios, considering working and weekend days (WE label in the table). Focusing on the Abilene case, E^{LC} is lower for weekend days than working

days since traffic is lower and hence more line cards are powered off. For example, considering the GA, E^{LC} is 821.75 kWh on 2004/09/02, and 689.5 kWh on 2004/08/29. Then, E^{LC} is higher for LFA: this is due to the fact that LFA switches off entire logical links rather than single lightpaths. However, turning off a logical link requires rerouting of the entire amount of traffic flowing on it, which is not always possible. This in turn leads to an increase of energy consumption, with E^{LC} always higher than 1645 kWh. On the contrary, both EWA and GA guarantee lower energy consumption, since they work on single lightpaths rather than logical links.

We then extend our analysis to the Germany17 and Géant networks. In particular, all the heuristics consume a larger amount of energy in Germany17 network than in Abilene. This is quite intuitive, since Germany17 has more nodes than Abilene. However, despite the Géant network being bigger than Germany17, the energy consumption on at least three out of four considered days in the Géant network is lower than in Germany17 on 2005/02/15. This can be explained by the traffic patterns on the days considered for evaluation of the heuristics, and by the overprovisioning of the SBNs. The overprovisioning of the Germany17 SBN is different from the overprovisioning of the other two networks despite using the same parameter γ . This is caused by the fact that the traffic data available for the Germany17 network contains only one day of traffic matrices with fine granularity (see Section 2.5 and [4]). Therefore the traffic matrices of granularity of one month were used to design the SBN (see Table 5.1). Since the values in the traffic matrices are average values over the considered period duration, the traffic peaks are averaged out, and are not included in D_{SBN} , which is calculated by taking maximum values for each demand over all traffic matrices (see Eq. (5.13)). Interestingly, EWA and GA consume similar amount of energy in the Germany17 network compared to LFA suggesting that, for this network, it is difficult to turn off single lightpaths. This is caused by the fact that traffic in the Germany17 network is centralized in Frankfurt (see Fig. 4.2(a)). Only 9 out of 97 logical links not attached to Frankfurt in the SBN are composed of more than one lightpath. The LFA tries to switch off each logical link, however, since most of the links are composed of just a single lightpath, there is little difference to EWA and GA, which try to switch off single lightpaths, and not logical links. The (relatively little) advantage in terms of energy consumption that EWA and GA have over LFA in the Germany17 network is caused by the logical links attached to Frankfurt (21 out of 26 are composed of more than one lightpath), and the possibility of establishing logical links not existing in the SBN. Finally, we point out that EWA guarantees lower energy consumption E^{LC} with respect to GA and LFA in the Géant network (Table 5.2).

In the following, we compare the heuristics considering the total energy consumption E^{TOT} . Table 5.3 reports the main results. In particular, E^{TOT} is higher than E^{LC} , since we are accounting also for the power of LCSs and FCSs. However, both GA and EWA are able to reduce the total energy consumption with respect to LFA, especially during weekends. These results confirm our intuition that minimizing the power consumption of line cards is of great benefit also for reducing the total power consumption of the network, assuming the possibility to switch off and on the LCSs and FCSs installed in the SBN. There is no significant difference between the results achieved by EWA and GA in terms of E^{TOT} and E^{LC} . The fact that EWA achieves sometimes better and sometimes worse results than GA can be explained by the fact that both EWA and GA are heuristic approaches.

As the next step, we compare the heuristics in terms of the reconfiguration ratio ξ . Table 5.4 reports the results for the different scenarios. Intuitively, a ratio ξ approaching 1 indicates that the amount of reconfigured traffic is comparable with the total amount of traffic demands. This

Table 5.2: Line Card Daily Energy Consumption E^{LC} [kWh].

Scenario	LFA	EWA	GA
2004/08/27 Abilene	2205.37	693.37	771.25
2004/08/28 Abilene (WE)	1673	587.75	697.875
2004/08/29 Abilene (WE)	1645.37	607.50	689.5
2004/09/02 Abilene	2255.62	712.75	821.75
2005/02/15 Germany17	2378.75	1946	1739
2005/06/07 Géant	3625.75	1312	1510.38
2005/06/10 Géant	2077	1231	1312.5
2005/06/11 Géant (WE)	2206.37	1136.12	1244.13
2005/06/12 Géant (WE)	2004.12	1061.5	1209.13

Table 5.3: Total Daily Energy Consumption E^{TOT} [kWh].

Scenario	LFA	EWA	GA
2004/08/27 Abilene	4246.25	1534.33	1612.21
2004/08/28 Abilene (WE)	3385.65	1428.71	1538.84
2004/08/29 Abilene (WE)	3338.97	1448.46	1530.46
2004/09/02 Abilene	4330.29	1589.77	1683.75
2005/02/15 Germany17	5467.78	4729.9	4302.93
2005/06/07 Géant	7753.17	3391.11	3632.37
2005/06/10 Géant	4847.14	3246.36	3266.59
2005/06/11 Géant (WE)	5081.82	2997.75	3107.17
2005/06/12 Géant (WE)	4685.66	3024.23	3046.11

condition is not of benefit in a network, since it might have a negative impact on the QoS of users due to temporary service disruptions and network congestion as many devices are changing their operational state. Focusing on the Abilene network, ξ is typically larger than 0.66 for LFA. On the contrary, ξ is clearly reduced by the other heuristics, being not greater than 0.15 and 0.18 for EWA and GA, respectively. This means that targeting explicitly the reconfiguration cost has a positive feedback on ξ . Similar considerations hold for the Germany17 and Géant networks. Finally, observe also that ξ does not significantly change over the different days, i.e., between working days and weekend days. This is due to the fact that a larger amount of traffic exchanged in the network (denominator of Eq. (5.11)) implies also a larger amount of reconfigured traffic (numerator of Eq. (5.11)).

Table 5.4: Reconfiguration Ratio ξ .

Scenario	LFA	EWA	GA
2004/08/27 Abilene	0.75	0.13	0.16
2004/08/28 Abilene (WE)	0.72	0.15	0.18
2004/08/29 Abilene (WE)	0.70	0.15	0.17
2004/09/02 Abilene	0.66	0.12	0.14
2005/02/15 Germany17	0.35	0.1	0.12
2005/06/07 Géant	0.40	0.09	0.11
2005/06/10 Géant	0.49	0.08	0.12
2005/06/11 Géant (WE)	0.48	0.1	0.13
2005/06/12 Géant (WE)	0.44	0.08	0.11

Table 5.5: Overload Ratio ϕ .

Scenario	LFA	EWA	GA
2004/08/27 Abilene	7×10^{-5}	0	0
2004/08/28 Abilene (WE)	0	0	0
2004/08/29 Abilene (WE)	3×10^{-5}	0	0
2004/09/02 Abilene	2×10^{-3}	0	0
2005/02/15 Germany17	5×10^{-3}	2×10^{-2}	4×10^{-6}
2005/06/07 Géant	8×10^{-4}	0	0
2005/06/10 Géant	0	0	0
2005/06/11 Géant (WE)	7×10^{-4}	0	0
2005/06/12 Géant (WE)	3×10^{-5}	3×10^{-5}	0

Finally, Table 5.5 reports the overload ratio over the different scenarios. Intuitively, ϕ should be kept as low as possible to limit the negative effects of packet dropping. Focusing on the Abilene network, $\phi = 0$ for both EWA and GA. On the contrary, $\phi > 0$ for different days when LFA is adopted. This is due to the intrinsic behavior of the algorithm: since LFA starts from SBN and tries to switch off logical links, in some cases the set of logical links provided by the SBN is not sufficient to satisfy the peak traffic.¹ Consequently, the overload is larger than zero, even if it is kept quite low, i.e., typically lower than 10^{-3} . On the contrary, both EWA and GA can accommodate the traffic demands more wisely, since they integrate the possibility to add lightpaths while satisfying the constraint of installed line cards in each network node. Table 5.5 reports also the results for the Germany17 network. In this case, $\phi > 0$ for all the heuristics, suggesting that the SBN is not sufficient to satisfy the traffic without exceeding the number of installed line cards. This is due to the fact that the SBN is designed with d_{SBN}^{ab} calculated using monthly averages of demand values (see Section 5.5). Finally, the Table 5.5 reports the results for the Géant network. In this case, $\phi < 8 \times 10^{-4}$ and $\phi < 3 \times 10^{-5}$ for LFA and EWA, respectively. However, the overload is limited, since it occurs only in the first 15-minute period on 2005/06/12, and can be avoided by using the network configuration from the last period of 2005/06/11 instead of the SBN as a starting point. GA avoids overload for all the periods in the Géant network.

5.6.1 Time variation

We then investigate the temporal behavior of the algorithms. In particular, we start considering the Abilene network² and the day 2004/08/27. Fig. 5.2(a) shows the power consumption of line cards versus time $P^{LC}(t)$. Power follows a day-night trend of traffic (Fig. 5.1(a)) for GA and EWA. Higher power is consumed by LFA, whose trend presents also spikes. In some cases LFA finds an aggressive configuration with many logical links switched off. Hence, the total power consumption is close to the other heuristics. However, the aggressive configuration is not able to accommodate traffic for a long period, i.e., more than one traffic matrix. This in turn leads to a new, less aggressive configuration of logical links with consistently higher power consumption. On the contrary, observe how GA and EWA are able to reduce power, with smooth power transitions. Note that the power consumption of all line cards installed in the SBN is more than 142 kW (not

¹Recall that we have adopted different time periods for designing the SBN and for evaluating the heuristics.

²We do not change the time zone of the original traffic matrices [4], therefore the day-night patterns in Abilene look shifted in time.

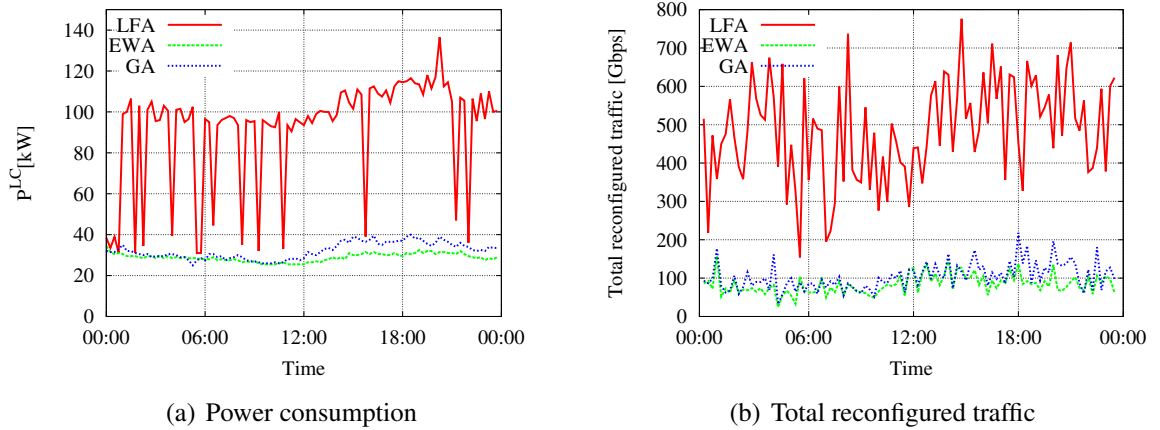


Figure 5.2: Power consumption of active line cards and total reconfigured traffic in the Abilene network on 2004/08/27.

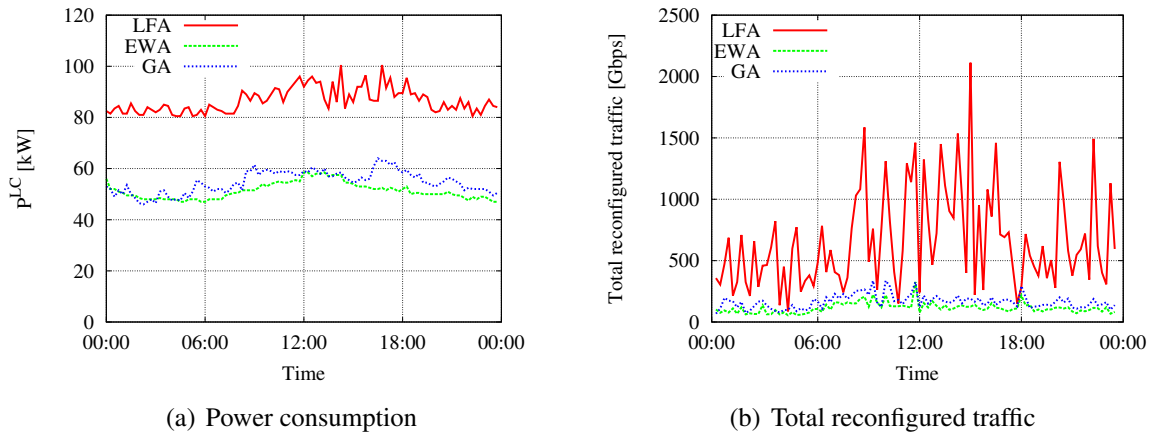


Figure 5.3: Power consumption of active line cards and total reconfigured traffic in the Géant network on 2005/06/10.

reported in the figure), thus we can conclude that our heuristics are effective in reducing the power consumption. To give more insight, Fig. 5.2(b) reports the total reconfigured traffic (sum of $r_{ij}^{ab}(t)$ in the fitness function (5.1)), expressed in Gbps. Both GA and EWA reduce the amount of traffic that is reconfigured, so that the reconfigured traffic is always smaller than 220 Gbps. On the contrary, LFA always requires higher amount of reconfigured traffic.

We then extend our analysis considering the Géant network and the day 2005/06/10. Fig. 5.3(a) reports power consumption of line cards versus time $P^{LC}(t)$. In this case, the power consumption of the always on network (SBN) is 276 kW, while our heuristics guarantee that power consumption is always lower than 105 kW. Again, a day-night behavior clearly emerges. Finally, the total reconfigured traffic is reported in Fig. 5.3(b). In this case, lower reconfigurations occur during the night for the three algorithms, suggesting that during the night the set of powered on devices does not frequently vary.

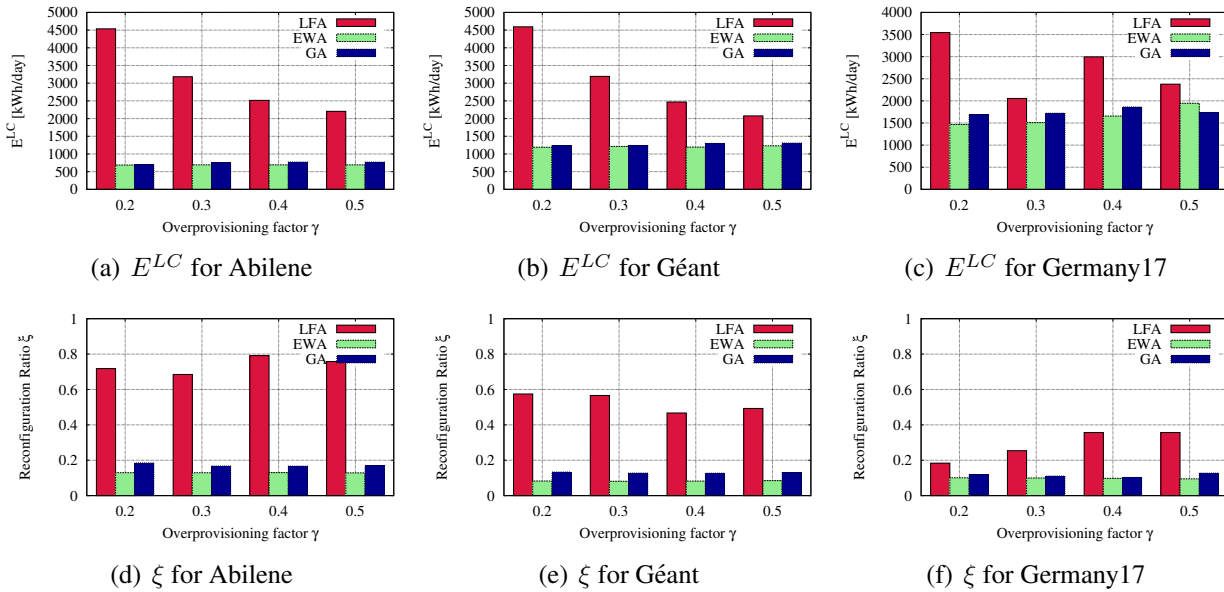


Figure 5.4: SBN Variation: Energy consumption and reconfiguration ratio in the Abilene and Géant networks on 2004/08/27 and 2005/06/10, respectively.

5.6.2 Impact of Static Base Network

In this section, we investigate the impact of the SBN design on the performance of our heuristics. In particular, we vary the overprovisioning factor γ in the range [0.2-0.5], trading between large and medium overprovisioning of SBNs. Intuitively, large overprovisioning requires more capacity to be installed, implying more power consumption but also more freedom in choosing which devices to power off. We assume traffic per node $d_{|V|}$ equal to 300 Gpn, and focus on the days 2004/08/27 for Abilene, 2005/06/10 for Géant and 2005/02/15 for Germany17.

Fig. 5.4 reports E^{LC} and ξ for the different scenarios. γ impacts the performance of LFA. In particular, for the Abilene scenario E^{LC} of LFA passes from 2200 kWh with $\gamma = 0.5$ to more than 4500 kWh with $\gamma = 0.2$. The reason why energy consumed by line cards in a network using LFA increases with increasing overprovisioning of the SBN is that the logical links that LFA did not switch off are still overprovisioned proportionally to the overprovisioning of the SBN. On the contrary, when considering the reconfiguration ratio ξ , the variation of SBN does not impact consistently the results. Moreover, when considering the Germany17 network, all the algorithms consume a similar amount of energy for $\gamma = 0.5$ and $\gamma = 0.3$. This is due to the fact that, for this type of network, it is more difficult to turn off line cards due to different behavior of traffic used to design the SBN. Differently from LFA, both EWA and GA present just minor oscillations of energy consumption and reconfiguration ratio for all the values of γ . This is due to the fact that both algorithms are not restricted to logical links existing in the SBN, and that they consider single lightpaths with a logical link, and not a logical link as a whole.




Figure 5.5: Load Variation: Energy consumption and reconfiguration ratio in the Abilene network on 2004/08/27.




Figure 5.6: Load Variation: Energy consumption and reconfiguration ratio in the Géant network on 2005/06/10.

5.6.3 Impact of load variation

Keeping the day-night traffic variation, we vary the total traffic per node $d_{|V|}$ between 100 Gpn and 500 Gpn in the SBN design phase with γ set to 0.5. We consider the day 2004/08/27 for Abilene, and the day 2005/06/10 for Géant. Fig. 5.5 reports the results for the Abilene network. E^{LC} is increasing with $d_{|V|}$ for all the algorithms (as expected). Interestingly, all the algorithms consume a similar amount of energy for $d_{|V|} = 100$ Gpn, since under little load a minimum number of line cards has to be always powered on to guarantee connectivity. However, as traffic increases, the LFA consumes more energy than GA and EWA. Moreover, ξ decreases as $d_{|V|}$ increases for EWA and GA, suggesting that these algorithms can wisely accommodate the increase of traffic by limiting the amount of reconfigurations.

Similar considerations hold also for the Géant network, reported in Fig. 5.6. It is interesting to note that the GA consumes slightly more energy than LFA under low load ($d_{|V|} = 100$ Gpn).

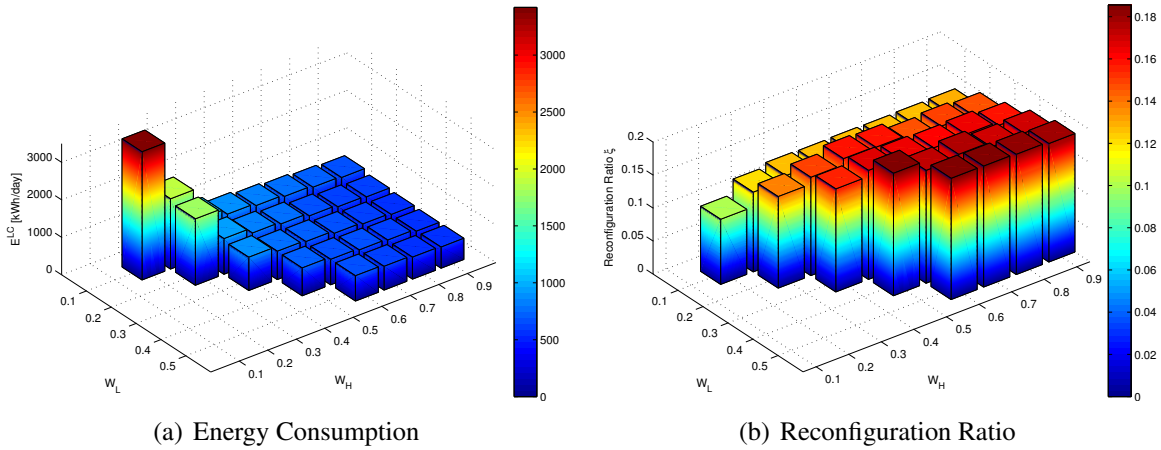


Figure 5.7: EWA Sensitivity Analysis for the Abilene network on 2004/08/27.

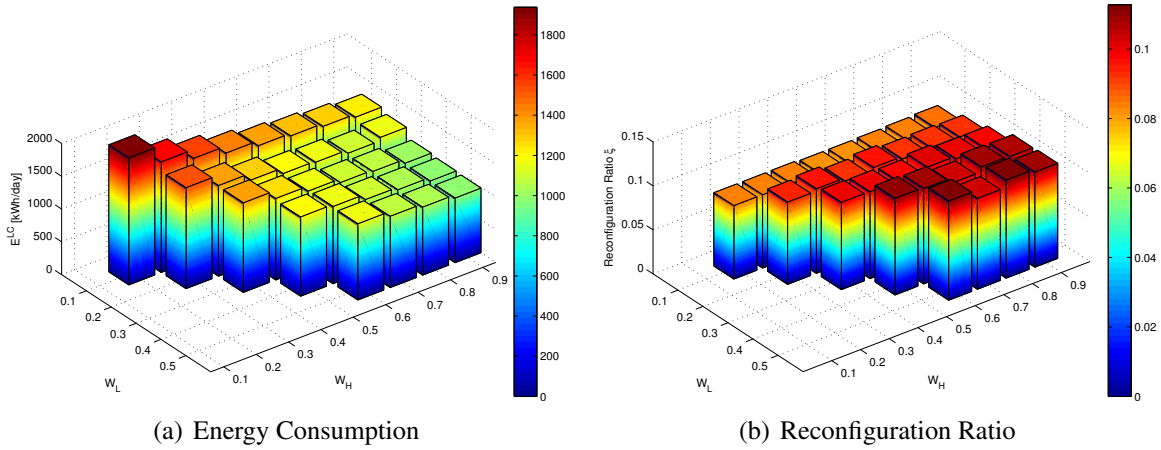


Figure 5.8: EWA Sensitivity Analysis for the Géant network on 2005/06/10.

5.6.4 Sensitivity analysis

In this section we evaluate the impact of the EWA parameters on E^{LC} and ξ . We adopt the following scenarios: the day 2004/08/27 for Abilene, and the day 2005/06/10 for Géant. We set $\gamma = 0.5$ and $d_{|V|} = 300$ Gpn.

We vary both the low threshold W_L and the high threshold W_H . E^{LC} and ξ for the Abilene network are reported in Fig. 5.7(a) and Fig. 5.7(b), respectively. Interestingly, the energy consumption steadily increases for $W_L \leq 0.3$ and $W_H \leq 0.3$, while the reconfiguration ratio is minimized. This is due to the fact that with this setting the algorithm tries to switch off lightpaths (low W_L) less frequently, but many devices are required to be powered on (overprovisioning due to low W_H). Conversely, with increasing values of watermarks ($W_L > 0.3$ and $W_H > 0.4$) the energy consumption decreases, but the reconfiguration ratio increases. Intuitively, for high values of W_L the algorithm is more aggressive in switching off devices, while for high W_H the algorithm tends to save more power since lightpath utilization increases. However, a good trade-off is guaranteed by

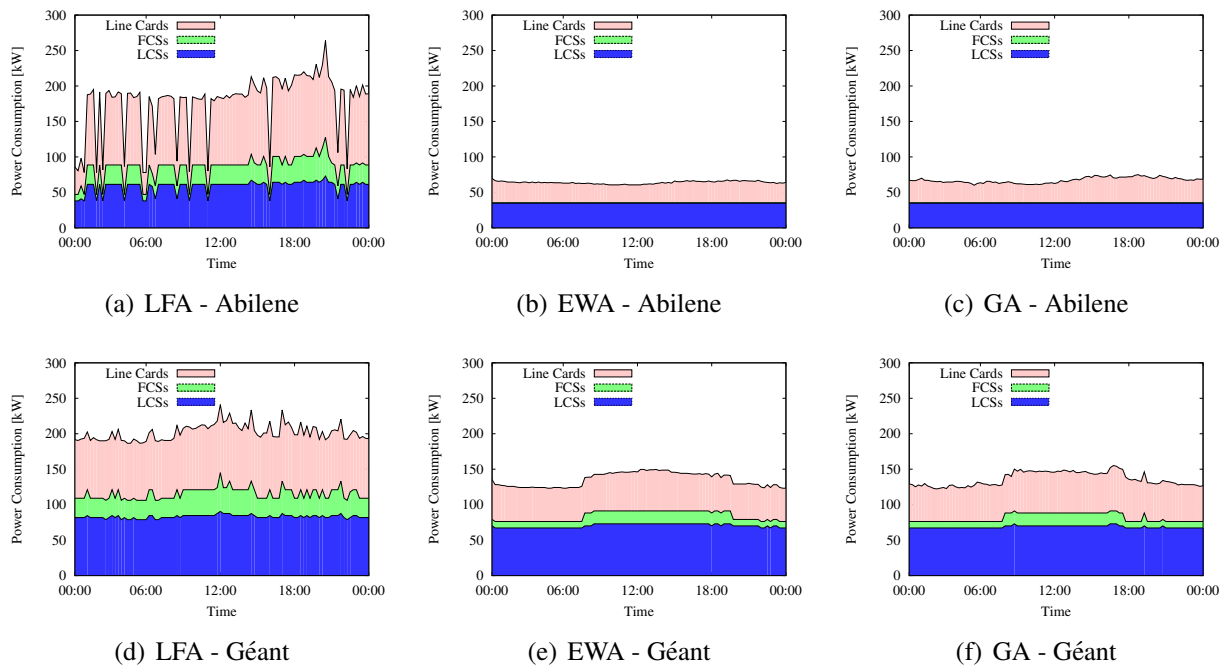


Figure 5.9: Power Breakdown for the Abilene and Géant networks on 2004/08/27 and 2005/06/10, respectively.

setting $W_L = 0.1$ and $W_H > 0.4$. Fig. 5.8 reports the results for the Géant network. Also here the energy is clearly minimized for high values of W_L and W_H (conversely for the reconfiguration ratio). A good trade-off in this case is to set $W_L = 0.1$ and $W_H = 0.9$.

5.6.5 Power breakdown

We then investigate how much energy is consumed by components of each type, differentiating between line cards, LCSs, and FCSs. In particular, the minimization of power due to line cards is targeted by the heuristics, while the power due to LCSs and FCSs is computed in a postprocessing phase. Fig. 5.9 on the top reports the power consumption versus time for the Abilene over 2004/08/27. The LFA heuristic (Fig. 5.9(a)) requires all types of components to be powered on over the whole day, with a large amount of power consumed by line cards during the day. The EWA heuristic (Fig. 5.9(b)) instead does not require any FCS, and a lower amount of power is consumed by line cards compared to LFA. Moreover, the power consumed by LCSs is constant, suggesting that EWA is able to find a stable solution for this type of devices. Finally, the power variation of GA, reported in Fig. 5.9(c) is similar to EWA, with higher variation of power consumed by line cards.

We then consider the Géant network over 2005/06/10 (the second row in Fig. 5.9). Interestingly, all the algorithms require to use FCSs. This is due to the fact that in this case, differently from Abilene, the distribution of traffic imposes to use more hops on average, and some nodes route a large amount of traffic. Moreover, both line cards and FCSs present a day-night trend, while the power consumption of LCSs is almost constant. This is however influenced by significantly higher power consumption of a FCS than the power of a LCS (9100 W and 2920 W according to

Table 5.6: Computation Times and Complexity.

(a) Computation Times [s]					(b) Worst-case Complexity	
Scenario	Statistic	LFA	EWA	GA	Algorithm	Worst-case complexity
Abilene 2004/08/27	median	0.84	0.01	25.8	LFA	$\mathcal{O}(L \log L + L \cdot V ^2)$
	average	0.85	0.03	26	EWA	$\mathcal{O}(V ^4 \log V + L Q^2 V ^2 \log V + G L \log L)$
	maximum	0.92	1.75	31	GA	$\mathcal{O}(M \cdot \kappa \cdot S \cdot V ^2)$
Géant 2005/06/10	median	0.89	0.13	147.57		
	average	0.89	0.31	144		
	maximum	0.98	15.06	225		
Germany17 2005/02/15	median	0.86	0.35	63.43		
	average	0.87	0.49	48		
	maximum	0.95	5.85	156		

Section 4.3.2), and significantly higher number of installed line cards than LCSs.

5.6.6 Computation times and complexity

Table 5.6(a) reports the comparison of the computation times required by the algorithms. EWA was simulated on a slower computer with different simulation framework than LFA and GA, therefore the results in Table 5.6(a) should be treated just as indicative ones.

LFA and EWA require the lowest amount of time to compute a solution. While the average computation times are close to median for LFA, EWA needs more time in the first period when switching off many lightpaths from the SBN. The adaptive behavior of EWA results in lower computation times in the next periods, which is reflected by different average and median computation values. GA requires more time than LFA and EWA, since several iterations are required to return the solution.

Focusing on computational complexity, the LFA first requires a sorting of the set of logical links, resulting in a time $\mathcal{O}(|L| \log |L|)$. Then, the computation of the paths and the constraints checks, which correspond to the flow conservation of the traffic demands and to the bandwidth utilization of the logical links can be done in time $\mathcal{O}(|V|^2)$. The procedure has to be repeated over the set of links, resulting in $\mathcal{O}(|L| \cdot |V|^2)$. Thus, the total algorithm complexity is $\mathcal{O}(|L| \log |L| + |L| \cdot |V|^2)$.

The computational complexity of GA is mainly dependent on the size of the network. Indeed, the complexity of the operations concerning the evolutionary process is dependent on the size of an individual which is equal to $|V|^2$. Thus, their complexity is $\mathcal{O}(|V|^2)$. In addition, the complexity for the operations of constraints checks, which are the same of LFA plus the check of the number of installed line cards in the SBN is also equal to $\mathcal{O}(|V|^2)$. These operations have to be repeated over all the individuals in the worst-case, resulting in a complexity of $\mathcal{O}(S \cdot |V|^2)$. Moreover, the complexity of the GA depends also on the size of the offspring κ and the maximum number of generations without improvement of the fitness function M . In particular, the total complexity is equal to $\mathcal{O}(M \cdot \kappa \cdot S \cdot |V|^2)$. The parameters M , κ and S constrain the search space of the GA. Consequently, these parameters determine the convergence time and the quality of the results. For instance, for low M , the amount of time employed by the GA is low, but it is likely to have worse result at the end, since few iterations are performed. Similarly, if a large number of individuals is generated at each generation, more solutions are examined and, thus, the search space increases

together with the convergence time. In this work, we keep M , κ and S fixed for all the simulations and we choose to use values that allow to find better solutions, at the cost of a larger execution time of the heuristic.

Finally, complexity of EWA is dependent on the input parameters, in particular on the network configuration in the previous time period. The IP routing of traffic demands $f_{ij}^{ab}(t-1)$ together with the current traffic matrix $D(t)$ influence the logical links where the utilization of the last lightpath violates the watermarks W_L and W_H . This in turn influences the attempts to establish and release lightpaths. The lightpath establishment is also dependent on the number of installed line cards in the SBN X_i^{LC} , and the number of established lightpaths forming logical links $y_l(t-1)$ influences the number of attempts to release lightpaths. We perform the worst-case complexity assessment of EWA assuming first that no demand is routable at the beginning of the algorithm, next that watermarks are violated on all logical links, and eventually that a new lightpath between each pair of line cards can be established.

We assess each step of the EWA presented in Alg. 1. Ensuring routability of demands over the network requires time $\mathcal{O}(|V|^2 \cdot |V|^2 \log |V|)$, since there are up to $|V \times V|$ traffic demands, and establishing of a new logical link for any demand may influence the routability of all demands, which need to be sorted ($|V|^2 \log |V|$). Sorting of the logical links according to the utilization of the last lightpath takes $\mathcal{O}(|L| \log |L|)$. Establishing of necessary lightpaths requires three types of operations: i) iterations over all the logical links where the high watermark W_H is exceeded ($|L|$), ii) sorting of the demands flowing through the logical link ($|V|^2 \log |V|$), iii) possible iterations over all demands flowing through the logical link ($|V|^2$). Finally the whole procedure has to be repeated in case a new lightpath is established between any line cards with the total number of line cards installed in the network defined as $Q = \sum_{i \in V} X_i^{LC}$. The overall complexity of establishing of the necessary lightpaths (consisting of the above three types of operations) is $\mathcal{O}(|L| \cdot (|V|^2 \log(|V|) + |V|^2) \cdot Q^2)$ (see [180]). Releasing unnecessary lightpaths takes $\mathcal{O}(G \cdot |L| \log |L|)$, where $G = \sum_{l \in L} y_l(t-1)$ denotes the number of established lightpaths in the previous time period, and $|L| \log |L|$ accounts for the fact that releasing each single lightpath forces sorting of logical links. The overall complexity of EWA (Alg. 1) can be therefore estimated in the worst-case as $\mathcal{O}(|V|^4 \log |V| + |L| Q^2 |V|^2 \log |V| + G |L| \log |L|)$.

The overall worst-case complexity of all algorithms is summarized in Table 5.6(b). All algorithms present complexity that scales at least as $|V|^2$. Moreover, complexity of both GA and EWA depends on the input parameters, which can be used to trade between quality of results and computation time. Finally, complexity of LFA is independent from the input parameters, and it is fixed given the SBN.

5.6.7 Monetary savings

Finally, we have estimated the impact of our heuristics on the monetary cost needed to power the network. In particular, we have selected the day 2004/08/27 for Abilene, the day 2005/06/10 for Géant, and the day 2005/02/15 for Germany17. By assuming that the same daily traffic profile is repeated over the whole year we have computed the total energy consumed by the network in a year, and the corresponding electricity costs.³ A cost of 0.0936 Euro per kWh is assumed, based on current statistics of electricity prices in Europe [181]. Table 5.7 reports the costs for the three

³Weekend days may further decrease the electricity consumption.

Table 5.7: Yearly Monetary Cost [thousands of Euros].

	Abilene	Géant	Germany17
ALL ON	280	526	260
LFA	145 (-48%)	166 (-69%)	186 (-28%)
EWA	52 (-81%)	110 (-79%)	161 (-38%)
GA	55 (-80%)	112 (-79%)	147 (-43%)

algorithms. Although not being the main focus of this paper, percentage savings with respect to an always on solution are reported in the parenthesis. Considering the Abilene network, an always on solution requires 280000 Euro per year. With LFA this amount can be almost halved, while with EWA and GA savings larger than 80% can be achieved, with 52000 Euro required to run the network with EWA. Higher savings can be achieved for the Géant network, with the lowest amount of electricity required by the EWA algorithm. Finally, savings between 28% and 43% are achieved with the Germany17 network. These results corroborate our intuition that reducing power consumption is of significant benefit also for reducing the associated monetary cost.

5.7 Conclusion

There are numerous challenges that the EA-ARSs need to handle. The proposed heuristic EWA takes into account most of the evaluation criteria identified in this chapter. It minds the constraint on the installed devices, impact on QoS and reconfiguration costs. Measured computation times are promisingly low (average below 0.5 s and median below 0.35 s). Clear triggering events are expected to be sufficient for collection of the required global network knowledge and centralized operation.

Our evaluation study was performed on several network scenarios taking as realistic parameters as possible. Future traffic was assumed to be unknown during the SBN design. The results indicate that switching off line cards is of great benefit to reduce the power consumption of the network. EWA performs better than LFA in terms of power consumption and reconfiguration costs. It achieves comparable results to the computationally intensive GA. Moreover, both GA and EWA are able to wisely trade between power and reconfigured traffic. This trade-off can be controlled by a proper setting of the input parameters. Traffic load impacts the power consumption of the network, but only marginally the associated reconfiguration ratio. Additionally, also the overprovisioning of the SBN has a limited impact on the reconfiguration ratio. Finally, we have shown that reducing power consumption is of significant benefit for reducing the associated monetary cost.

The main open issues for EWA are the considerations of physical layer constraints and of protection. Availability of traffic data of a certain granularity Δt limits the possibility to fairly evaluate EA-ARSs under the assumption of unknown traffic. Evaluations on experimental testbeds are therefore required to verify implementation issues.

6 Implementability

A natural way towards a common use of an energy saving approach on real networks is its verification on a testbed or even on a real network. The goal of this chapter is to provide an outlook of implementation issues for Energy-Aware Adaptive Routing Solutions (EA-ARSs) and to verify whether the EA-ARSs discussed in this thesis are feasible in real world. A successful execution of an EA-ARS itself is insufficient to save energy. We outline implementation issues that need to be handled in order to make energy saving with EA-ARSs a reality (future work). An overview of the work related to experimental activities is provided. Based on the related work, steps needed for energy saving are identified, and possible methods of implementation of EA-ARSs in an Internet Protocol (IP)-over-Wavelength Division Multiplexing (WDM) network are discussed. Eventually, the approaches Fixed Upper Fixed Lower (FUFL) and Dynamic Upper Fixed Lower (DUFL) are demonstrated on an IP-over-Gigabit Ethernet (GbE) testbed. No IP-over-WDM testbed was available to us, however the same control and management mechanisms can be used for the IP-over-WDM and IP-over-GbE networks. We focus on Generalized MultiProtocol Label Switching (GMPLS). Therefore some relevant GMPLS definitions are discussed first (extending Chapter 2). Parts of Chapter 6 have been published in [182].

6.1 Definitions required for the management and control with GMPLS

We assume GMPLS to be used for the management and control of the IP-over-WDM networks. We introduce some definitions related to GMPLS based on [24]. GMPLS is concerned with merging the concepts of data forwarding in different types of networks (packet, frame, cell) and the concept of transport networks (WDM, Time Division Multiplexing (TDM)). GMPLS-capable routers are referred to as Label Switching Routers (LSRs).

Network resources available for traffic engineering path computation are modeled in GMPLS as Traffic Engineering (TE) links. Two types of TE links are considered in this thesis. First, the physical fiber link, and second, the logical link. A Traffic Engineering Database (TED) stores information about all TE links. The most relevant information describing TE links in the context of green networking is the traffic engineering metric, the maximum link bandwidth and the maximum reservable bandwidth (see Table 5.1 of [24] for a complete list). TED is used by a Path Computation Element (PCE) to compute the traffic engineered paths (lightpaths and traffic route from its source node to its target node) using a specific algorithm (e.g., Dijkstra).

PCE is separated from the routing protocol (e.g., Open Shortest Path First-Traffic Engineering (OSPF-TE)). It calculates a path through the network (Label Switched Path (LSP)) using information stored in TED. The path is stored in an Explicit Route Object (ERO) (a data structure containing a sequentially ordered list of LSRs).

OSPF-TE is a link state routing protocol used in GMPLS. It is an extension of the traditional Open Shortest Path First (OSPF) link state routing protocol, where Link State Advertisements (LSAs) flooded through the network (Traffic Engineering Link State Advertisements (TE-LSAs)) contain information about router addresses or TE links.

Link State DataBase (LSDB) is a data structure within any given link state routing protocol speaker. The LSDB is a repository for any LSAs in a particular routing area. A single routing area is assumed in this work.

The signaling protocol Resource reSerVation Protocol-Traffic Engineering (RSVP-TE)¹ is used to instantiate LSPs, where labels refer to packet or wavelength in this work. The following messages are used by RSVP-TE: Path (LSP Setup), Resv (LSP Accept), ResvConfirm (LSP Confirm), PathErr (LSP Upstream error), ResvErr (LSP Downstream error), PathTear (LSP Downstream release), PathErr with a flag indicating “Path state removed” (LSP Upstream release), and Notify (LSP Notify).

Network Management System (NMS) is a central management station or application that has a view of the whole network, and can control and configure all of the devices in the network. It uses proprietary management interfaces (either Command Line Interface (CLI) or Graphical User Interface (GUI) specific for each device) and usually a standardized management protocol such as Simple Network Management Protocol (SNMP). Network devices can be accessed locally (e.g., via RS232 interface) or via web. The management messages are carried by a management network. The connectivity is usually provided by an IP management network. However, if there is in-band or in-fiber control plane communication between network devices, then the management messages may also be carried in this way. The NMS (and also Operations Support System (OSS)) has to be able to manage all network devices via Element Management Systems (EMSs). The Management Interface Base (MIB) is a global distributed database for management and control of SNMP-capable devices. A MIB module is a collection of individual objects and tables, each of which contains a value that describes configuration or status of a manageable entity or logical entity of the same type. We present the MIB modules relevant to the EA-ARs assuming presence of the corresponding tables (e.g., Routing Information Base (RIB)) in the routers. The details of its implementations on the routers commercially available today is unknown, even when looking at the product manuals (see e.g., the Interfaces and Routing Configuration Guide of Junos Software [183]). Software Defined Networking (SDN) is a promising approach for green networking, however we leave it for future work in this thesis.

The following MIB modules are particularly relevant for green networking, since they are responsible for links (their presence or not) and (re)routing.

1. GMPLS LSR MIB module (GMPLS-LSR-STD-MIB, which extends the MPLS-LSR-STD-MIB) – it contains four tables originating from MultiProtocol Label Switching (MPLS): a table of GMPLS capable interfaces (gmplsInterfaceTable), a table of “in-segments” and a table of “out-segments” describing the downstream legs and upstream legs of LSPs respectively (gmplsInSegmentTable and gmplsOutSegmentTable), as well as a table of “cross-connects” that shows the relationship between in- and out-segments (mplsXCTable, no extension in GMPLS). Indication of LSPs which use TE links in sleep mode can be stored in extra_Parameters_Table related to gmplsInSegmentTable and gmplsOutSegmentTable.

¹We leave the Constraint-based Routed Label Distribution Protocol (CR-LDP) aside in this thesis.

2. GMPLS TE MIB module (GMPLS-TE-STD-MIB, which extends the MPLS-TE-STD-MIB) – it models and controls GMPLS TE LSPs. The following tables are particularly relevant from the perspective of green networking: `gmplsTunnelTable` containing active and configured LSPs; `mplsTunnelResourceTable`, where resource requirements and usage of each LSP are stored; and `gmplsTunnelHopTable`, `gmplsTunnelCHopTable` and `gmplsTunnelARHopTable` for specification, computation and recording of the path taken by the LSP, respectively.
3. TE Link MIB module (TE-LINK-STD-MIB), which is responsible for setting up and using link bundles configured from TE links. Its `teLinkTable` contains entries representing TE links (including bundled links) and their generic TE parameters.

6.2 Experimental activities – a survey

There is a limited amount of related work dealing with implementation of energy saving mechanisms through selective on/off switching of network elements during periods of low load. To the best of our knowledge, it is basically limited to two activities, i.e., experiments with the MiDORi [Multi-(layer, path and resources) Dynamically Optimized Routing] Network Technologies project [184] and experiments on the CARISMA testbed. The main general difference between them is that new network devices are developed with MiDORi, while off-the-shelf equipment of higher capacities is used within CARISMA.

6.2.1 MiDORi

The focus of the MiDORi is set on energy saving in the GbE network by selectively powering off network interfaces under hop-limit and bandwidth constraints. Powering on/off of the whole transit routers or their parts is also considered. Starting with [149], the authors propose a solution which effectively creates all on/off combinations of links in the network. The solution that can carry the whole traffic and consumes the lowest power (considering moving some virtual routers in order to power off also nodes) is selected. The Beeler's algorithm (see Section 3.2.3) is compared with the proposed any-order pattern algorithm, which additionally guarantees maximum hop number of the path and disjoint multi-route link divergence for reliable communications.

Both algorithms have been implemented on a parallel reconfigurable processor DAPDNA-2 processor (by IPFlex Inc.), where DAP stands for Digital Application Processor and DNA stands for Distributed Network Architecture. The algorithms are compared in terms of computation time using the DAPDNA-2 processor (referred to as proposal) and a Pentium 4 processor (referred to as conventional) based on experiments and theoretical estimation. The DAPDNA-2 implementation turns out to be 40 times faster.

While the algorithms are centrally executed by the PCE, the authors consider also link on/off control protocols. Extensions to GMPLS are proposed. More precisely, a new state (power on/off) of physical and TE links is proposed for the OSPF. The controlling is supposed to be performed over an out of fiber channel. New commands “Link power up” and “Link power down” are proposed for RSVP-TE and (alternatively) Link Management Protocol (LMP). The LSPs are supposed to be added a new status flag “power off by MiDORi TE”. Details of the extensions are not available in [149].

Further publications related to MiDORi provide extensions of [149]. In particular, a prototype Layer 2 (L2) switch is introduced in [185]. A depth- d algorithm for searching the optimal configuration of the logical topology is proposed in [63], where d determines the maximum number of links that are attempted to be switched off. The depth- d algorithm makes sure that the maximum link utilization does not exceed 1.0. Simulation results in terms of reduced number of links versus load (for two values of d compared against optimum solution) are provided for the NSFNET network.

The following contributions are made in [65]. First, clear steps for the energy saving in the MiDORi architecture are described, i.e., (1) Traffic monitoring; (2) Calculation of energy-efficient logical topology by PCE; (3) Reconfiguration of the network. Second, experiments on a 6-node 7-link network using the depth- d algorithm were conducted, and results similar to the ones from [63] were presented together with the calculation times in the range $0.01\text{--}10^5$ s for networks with 10–100 nodes, and depth- d in the range 1–4. In [65] the authors report also total current of the prototype L2 switches (2.765–2.831 A), and mention Ethernet Virtual Local Area Networks (VLANs) as the way of controlling traffic paths.

More details of the GbE L2 switch are provided in the block diagram presented in [186]. The switch can count traffic of each LSP (VLAN) and each GbE link. The power consumption of the switch can be read via a command and a current meter. The presented switch has eight GbE links and is controlled remotely (power on/off state of each link and each fabric) using telnet via a linux based control card which is one of the few parts constantly powered up. The feature of Self Organizing Network (SON) is mentioned for the first time (marked as under development) in [186]. The authors demonstrated MiDORi on a fully meshed 6-node network testbed using the depth-1 algorithm with generic Quality of Service (QoS) restrictions. Six traffic generators/receivers were used, however the traffic assumptions were not detailed except for the fact that low traffic to high traffic ratio equals 1:5. The following steps are distinguished (extension from [65]): Step (1) Reading the traffic counters of each VLAN by the PCE and calculating average values; Step (2) Execution of the depth-1 algorithm at the PCE to obtain the logical topology and VLAN paths; Step (3-1) Powering on/off links in all switches (remotely by the PCE) according to the topology from Step (2); Step (3-2) Reconfiguration of the VLAN network topology according to the path calculation from Step (2).

Execution of the Steps (1)–(3-2) is repeated every X minutes, however the authors do not report X in [186]. Parallel and serial control of the switches were considered, with the parallel control taking significantly less time both during the traffic increase and traffic decrease (233.7–243.9 s vs. 61.8–68.7 s, for serial vs. parallel control respectively) The results show that the calculation of topology (Step (2)) takes marginal time (0.004–0.006 s). Duration of Steps (1), (3-2) and (3-3) takes 23.8–112.7 s in the serial control, and 7.2–29.0 s in the parallel control. About one third of the links can be powered off in the low-demand hour, however an inconsistent values for the total number of links is reported in Fig. 4 (30) and Fig. 7 (31) of [186].

The concept of SON is continued in [187] using the depth- d algorithm again. The authors point out that the MiDORi GMPLS supports multiple layers, multiple paths, and multiple resources. They explain again the OSPF extension (relation between physical links and TE links), LMP extension (power on/off control function using the LMP ChannelStatus message with Ack and IP Control Channel (IPCC) always up), and Resource reSerVation Protocol (RSVP) extension (Power control request in the ADMIN STATUS object for LSP status flag). Moreover, they provide a web link to the MiDORi GMPLS software available at [184]. The authors mention the 16-port

GbE switch, which they developed additionally to the 8-port switch presented in [186]. Demonstration on a 5-node 7-link network is performed showing that total switch power consumption can be reduced from 283.1 W to 276.1 W. The results of the reconfiguration times from [186] are also summarized in [187]. Additionally, the authors point out that their prototype switch does not support a “make-before-break” VLAN reconfiguration, and therefore data disruption occurs over the 29 s of VLAN reconfiguration. Eventually, the authors mention a MiDORi GMPLS optical switch, which they developed. It is also controlled via telnet by the PCE implemented on a small linux box. The optical switch allows the authors to demonstrate the multi-layer GMPLS signaling between a Lambda Switch Capable (LSC) layer and a Layer 2 Switch Capable (L2SC) layer. The variation of traffic takes place in 5-minute intervals.

In [188], the authors show the energy consumption (in Wh without specifying the considered time period and details of the traffic data) on a 4-node full mesh network. The energy saving reaches up to 23.8%.

Prim’s algorithm for Point-to-MultiPoint (P2MP) communication is proposed in [189]. It is compared with the depth-d algorithm modified for P2MP operation. The performance of the algorithm is evaluated in a simulative way on a 40-node, 90-link network. The capacity of each link is 1 Gbps, and the “required traffic” is reported as “15 Mbps/ static flow”. The number of multicast groups, site of occurrence of request, and the number of members in one multicast group are randomly assigned. The results show that the Prim’s algorithm outperforms the modified depth-d algorithm in terms of calculation time and the percentage of switched off links.

The experiments with the multi-layer network using GbE switches and the optical switch from [187] are continued in [190] using the extension of the GMPLS. Namely 4 Ethernet switches out of the considered 6 are connected to the optical switch. The demonstrated power saving (9.4 W corresponding to 6 ports) is low, but shows that the MiDORi network technology is potentially feasible in a high speed and power consuming interface located in a large scale network environment.

Eventually, the results are summarized in [64], which directly extends [149]. It includes the Beeler’s algorithm, the any-order pattern algorithm, their simulative evaluation on the NSFNET network loaded with uniformly generated inter-node traffic, the GMPLS extensions (OSPF, RSVP and LMP), the 8-port GbE switch development, and the same results as in [186] (fully meshed 6 nodes network testbed).

The GMPLS extensions developed within the MiDORi project have been proposed to the Internet Engineering Task Force (IETF) [191].

6.2.2 Experiments on the CARISMA testbed

Alcatel-Lucent Bell Labs (ALBF) jointly with Universitat Politècnica de Catalunya (UPC) proposed the extension of GMPLS in [21], and performed experiments on the CARISMA testbed. More specifically, the authors of [21] propose to introduce a new bit ‘S’ to the RSVP-TE Path and Resv messages. The bit ‘S’ used jointly with the already existing bit ‘A’ allows distinguishing between the following states of an Electrical-Optical (OE) device (such as colored line card, transponder or regenerator): Up, Idle, Down, Damaged.

The proposed extension has been evaluated in the CARISMA testbed available at UPC premises in Barcelona. The testbed was configured according to a Pan-European network composed of 16 nodes and 23 links with 10 bidirectional 100-Gbps-wavelengths per link. 20 add/drop transponders and 10 regenerators were used at each node. The testbed was loaded with uniformly distributed

lightpath requests according to the Poisson model with average holding time equal to 3 hours. The load is varied between 40 and 80 Erlangs.

Differentiated provisioning of connection requests is considered in [25] for gold, silver and best-effort traffic. In this case, the 'A' and 'S' bits of the ADMIN Status object of the RSVP-TE Path message are used in the following way: i) 'A'=0 and 'S'=0 indicate the OE devices in up state which must be used to allocate gold requests; ii) 'A'=0 and 'S'=1 indicate the OE devices in idle state which are needed to allocate silver requests; iii) 'A'=1 and 'S'=0 indicate the OE devices in down state which can be used to allocate best-effort requests. Differently to [21], availability of regenerators and wavelengths on links is disseminated over the network using the proposed extension of the GMPLS OSPF-TE protocol. A new sub-TLV (Type Length Value) (named TSP Status) is introduced in the OSPF-TE opaque LSAs containing the number of up, down and idle transponders in a node (a regenerator corresponds to two transponders). This sub-TLV is inserted into a Node Information top level TLV (type 5, see Fig. 1 of [25]). Using the OSPF-TE opaque LSAs, the PCE can populate its TED with wavelength and regenerator availability information, which is used for computation of end-to-end routes.

The same topology as in [21] is used for the experimental study on the CARISMA testbed. The service class distribution is divided as 20/30/50% for gold/silver/best-effort traffic, respectively. Different shares of resources are reserved for different classes of traffic. Pre-reservation of resources is implemented in the PCE to avoid contention of resources among different lightpaths under establishment. Results on the blocking ratio, number of OE devices in Up/Idle/Down states and power consumption per active LSP are reported.

The experimental activities on a testbed reported in [21, 25] were restricted to protocol information exchanges. Due to unavailability of transponders, the idle-up and down-up state transition times were assumed and not measured. The assumed transition times equal 20 ms and 60 s, respectively.

6.3 General steps needed for energy saving

Energy saving using EA-ARSs requires not only calculation of proper network configuration by an EA-ARS based on given input data (Chapter 5), but also other actions. Fig. 6.1 shows schematically the following four steps: (1) traffic monitoring; (2) validation of events triggering network reconfiguration; (3) calculation of the energy-efficient network configuration; (4) network reconfiguration.

The implementability is closely related to the operation of the EA-ARS (centralized or distributed) and the network knowledge (in general local or global) required by the EA-ARS. Distributed operation requires implementation of the EA-ARS in each network node, and delivery of the global network knowledge to it (if needed). In turn, the centralized operation requires implementation of the EA-ARS on one device, delivery of the global network knowledge to it (if needed), and eventually delivery of the computed solutions to all network nodes.

We assume a centralized implementation on a PCE (corresponding to the dashed box in Fig. 6.1) without losing the generality of the problem, since each distributed EA-ARS can be implemented in a centralized manner. Moreover, the centralized operation is more challenging on one hand (need for transportation of the network knowledge information, and the computed network configuration solution), and more universal on the other hand, since it allows to overcome the limitation of the




Figure 6.1: Steps needed for energy saving (traffic rates in Gbps, IP link capacities in number of lightpaths denoted as λ in the figure).

access to network devices such as IP routers being a black-box for the user (network operator) who can only execute given (limited) set of commands on it via CLI, GUI or web interface. To the best of our knowledge, there is no possibility to implement EA-ARSs on the off-the-shelf backbone routers available today.

6.4 Available methods

We look at the available methods for implementation of an EA-ARS such as FUFL in an IP-over-WDM network with the GMPLS control plane and centralized operation. The discussion below is based on [3, 64, 116].

6.4.1 Traffic monitoring and communication of the monitored information to the central PCE

IP routers have the possibility to monitor traffic on its interfaces [182]. It can be communicated by means of SNMP to the PCE using a direct connection between a PCE and IP routers via e.g., Fast Ethernet (FE). This approach can be implemented today using off-the-shelf equipment. However, a drawback of this approach is that the monitored information is communicated via the best-effort IP network together with the data traffic.

Another possibility to communicate the monitored traffic information is to use the out-of-bound signaling [64]. The out-of-bound signaling can be realized in the WDM layer by establishing a direct, dedicated lightpath between each IP router and the PCE. This approach is however expensive (also from the energy-efficiency point of view), since a dedicated lightpath and a pair of line cards is needed for each network device under control.

Eventually, extension of the OSPF can be used. Namely, the TE-LSAs proposed in [192] can be flooded in a similar way as regular LSAs. The link load can be calculated as a difference between the maximum bandwidth and unreserved bandwidth carried by the TE-LSAs.

Furthermore, a new sub-TLV named TSP Status (i.e., Transponder Status) is proposed for opaque LSAs in [25] to distribute the the number of up, down and idle transponders in a node.

6.4.2 Validation of triggering events

Two main types of triggering events can be distinguished: (1) expiry of timers, where calculation of network configuration is performed every x seconds (typically 5–15 minutes); (2) exceeding of load thresholds in certain points of the network [18, 176] or the total traffic in the network [106].

In the case (1), the validation is periodic, while in the case (2) it is not. The validation (2) itself is trivial, since it is just an execution of a simple “if greater than” test comparing the measured values reported with LSAs against predefined threshold values. Time needed to perform this test is negligible regardless of its implementation in any programming language.

6.4.3 Calculation of network configuration

There are many different approaches to calculate a new network configuration using an EA-ARS as elaborated in Section 3.2. Apart from the complexity of each EA-ARS itself, the implementation will influence calculation time of network configuration. An extreme case of implementation of an exhaustive search (Beeler’s algorithm) is tackled in [64].

As far as the centralized operation is assumed, the central PCE can be operated on a normal Personal Computer (PC), and therefore any programming or scripting language can be used to implement the EA-ARS.

Moving towards distributed implementation, the experimental activities can use platforms such as Field Programmable Gate Arrays (FPGAs) or software routers. Successful implementations on such platforms should encourage equipment vendors to implement EA-ARSs on their devices. Local implementation on commercially available devices is reserved to the equipment vendors, since no access to the firmware is given to the users (i.e., network operators).

The critical point for the implementation is to mind the time required to calculate the network configuration. Therefore solving of complex Mixed-Integer Linear Programming (MILP) problems is out of scope for the EA-ARSs.

6.4.4 Communication of new network configuration and the network reconfiguration

Once a new network configuration (different from the old one) is found, it needs to be communicated to the corresponding network nodes, so that the network reconfiguration process can start.

The process of communication is similar to the one presented in Section 6.4.1, i.e., via a direct connection between a PCE and IP routers (e.g., logging to the routers via telnet [186]), out-of-bound signaling, and GMPLS extensions.

We elaborate on the last method. First, as an extension of the OSPF, a TE-Metric is proposed [116, 192] by Zhang et al. The TE-Metric is flooded in the network as LSAs. Its value equal to the OSPF weight indicates a need to turn the link on. Its maximum value indicates that the link should be turned off.

Second, extensions to RSVP-TE are proposed. Usage of the ADMIN STATUS object of the RSVP-TE path/resv messages is proposed in [3, 187], where an extra bit [3] (flag [187]) is added

Table 6.1: Current (“A”) and proposed (“S”) bits in the signaling protocol (ADMIN STATUS object of the RSVP-TE path/resv messages) [3].

Power state	Bit “A”	Bit “S”
On	0	0
Sleep	0	1
Off	1	0
Damaged	1	1

to communicate the administratively setting-up or tearing down a specific LSP (lightpath). In other words, the added flag indicates whether the LSP should use slept link(s) or not. During the power-on (activation) procedure, the TE links (parts of a lightpath corresponding to a single fibre link) are powered on sequentially over the whole LSP. The power-off (deactivation) procedure assumes either deactivation of all TE links over the whole LSP (what requires rerouting other LSPs if they use the same TE links), or powering off all but not shared TE links. The proposed bit combinations (including the new bit “A”) in the signaling protocol for describing the device states is shown in Table 6.1 [3].

Third, similar extensions as for RSVP-TE are proposed also for LMP [187] by Okamoto et al. Changing of the state of a link (on/off) is performed by exchanging a request and ack between the corresponding routers (LSRs) and changing the TE link status. This can be done using the LMP ChannelStatus message [187], what requires a dedicated IPCC to be always up. The state of the links (TE links [64]) needs to be distinguished by the OSPF protocol, so that the links in the sleep state are taken into account.

The network reconfiguration itself needs to make sure that no traffic loss occurs. This can be ensured by the make-before-break mechanism. During reconfiguration:

- first new links (lighpaths) are established (which involves activation of network interfaces),
- then all the traffic is rerouted to the new paths,
- and eventually the unnecessary links (lightpaths) are released (which includes deactivation of network interfaces).

It must be ensured during the traffic rerouting that not only the new traffic takes new paths, but also the old traffic that is already on the way towards its target node on the old paths reaches its target. After traffic rerouting, a router should wait the time corresponding to propagation delay of the link before the link is deactivated. Similar care needs to be taken when activating a link. The router must make sure that both ends of a link are ready before performing the actual traffic rerouting.

The time of network reconfiguration is not so critical, because the new resources are allocated, before releasing the unnecessary ones. The time needed to switch on a device (until it becomes fully operational) varies depending on their type. According to [21], lasers at the optical emitter and receiver require approx. 60 s, and other devices require from micro- to mili-seconds to be operational.

The changes in the network are monitored by the central NMS. Therefore the relevant MIB modules (particularly GMPLS-LSR-STD-MIB, GMPLS-TE-STD-MIB, and TE-LINK-STD-MIB) need to be updated.

6.5 Implementation on a testbed

The proposed approaches FUFL and DUFL (Section 4.2) have been implemented by Edion Tego (Fondazione Ugo Bordoni) on an IP-over-GbE testbed [193] located at the Institute of Communications and Information Technology (Istituto Superiore delle Comunicazioni e delle Tecnologie dell'Informazione ISCOM) of the Italian Ministry of Economic Development. The IP-over-GbE testbed was the only one available to us, however the same control and management mechanisms can be used for the IP-over-WDM and the IP-over-GbE networks. For the sake of demonstration of the energy saving using the EA-ARs, optical GbE links are used instead of the WDM lightpaths. However, not only the process of activation/deactivation of network devices matters. Also the time needed to access devices needs to be taken into account, as well as serial or parallel operation of the reconfiguration on different devices in the network.

6.5.1 Testbed setting

A central PC is used for network management and control. It is connected directly over FE links to three IP routers available to us on the testbed. Monitored traffic information is communicated by means of SNMP. Triggering of calculation of a new network configuration based on the monitored traffic, as well as the calculation of the new network configuration itself are realized using bash scripting. Eventually, network reconfiguration is performed by logging via telnet into the IP routers, and executing commands to perform rerouting and activation/deactivation of GbE interfaces.

6.5.1.1 Base logical topology

We focus on the core part consisting of one Juniper M10 and two M10i IP routers interconnected by 50 km of fiber cable into a bidirectional physical ring. The logical topology is composed of the IP routers interconnected by GbE optical links. All parallel GbE optical links between a node pair form a logical link. The logical topology is controlled in a centralized way [194] from a linux PC.

We configure the testbed in order to obtain a simple scenario that allows demonstration of FUFL and DUFL. The base logical topology together with all traffic demands (source-target) is presented in Fig. 6.2. The nodes A–D and H–K represent traffic generators and sinks attached to nodes E and G. The part representing a core network consists of the nodes E, F and G interconnected by three logical links, each formed by two GbE optical links.

6.5.1.2 Traffic and routing

Traffic and its routing over the base logical topology have been chosen in an artificial manner so that FUFL and DUFL operation can be demonstrated on the available testbed. The Ethernet Testing Platform Spirent SPT-3U, Anritsu MD1230B and a linux PC have been used to generate and terminate traffic. Two types of traffic are used: (i) random traffic with specified minimum and

Table 6.2: Traffic demands (bidirectional) in the network.

Traffic Demand	Traffic Type	Min [Gbps]	Max [Gbps]	Period [s]
A–H	random	0.97	1	-
B–I	sine-like	0	0.1	200
C–J	sine-like	0	0.1	200
D–K	random	0.97	1	-

maximum values; (ii) sine-like traffic with specified minimum and maximum values, as well as period length in seconds. The sine-like traffic consists of halves of sine periods and of idle periods (as indicated in Fig. 6.5(c)), what is determined by the traffic generators.

The maximum value of the sine-like traffic that our traffic generators can produce is 100 Mbps. Corresponding value for the random traffic is 1000 Mbps. A summary of the traffic inserted into the network is provided in Table 6.2. Traffic is generated in both directions.

The IP routing in the base network indicated in Fig. 6.2 has been chosen so that all the logical links carry traffic, and that the load exceeds the capacity of a single GbE optical link. It gives us the opportunity to see what happens with the traffic on a logical link when the whole logical link or just one out of two parallel GbE links is switched off in the low-demand hour.

The inefficient utilization of the logical links is caused by the limitations of the traffic generators. This constitutes no obstacle for showing the operation of energy-saving approaches, but explicitly provides potential for switching off the GbE interfaces.

6.5.1.3 Power consumption

Power consumption of interfaces has been measured offline using Precision Power Analyzer N4L PPA2530 and a method similar to the one from [19]. Namely, power consumption of an IP router was measured twice, i.e., when a GbE interface was active and when it was inactive.² The sub-

²Further measurements of power consumption of an IP router with all interfaces physically removed showed a difference of less than 0.5 W with respect to the router with an inactive GbE interface installed.




Figure 6.2: Base logical topology.




Figure 6.3: Differentiation between T_M (time period over which the constantly monitored traffic is averaged) and T_L (the time period after which triggering events are validated).

traction of these two values determined the power consumption of the GbE interface (11.07 W for the M10 router and 8.9 W for the M10i router). Power consumption of the Juniper’s M10 and M10i routers with all interfaces shutdown equals (according to our measurements) 186.15 W and 112.5 W, respectively. The power consumption values reported above are lower than the ones reported in [19] for Cisco 7507, and in [187] for the MiDORi Ethernet switch.

6.5.2 Methods

We explain the methods used in each step needed for energy saving according to Fig. 6.1.

6.5.2.1 Traffic monitoring

The most intuitive approach to traffic monitoring is to perform the monitoring constantly. This approach is impossible in the simulative works performed so far, due to unavailability of input traffic data sets originating from measurements and covering traffic between all node pairs in the network over sufficiently long period of time with sufficient time granularity (IP packet arrival/departure level). An experimental activity overcomes this limitation. However, there is still some level of freedom in setting the traffic monitoring in a digital system such as a telecommunication network (testbed), namely the time period T_M over which the constantly monitored traffic data is averaged and provided for evaluation triggering calculation of a new network configuration.

6.5.2.2 Validation of triggering events

The second level of freedom determines the frequency at which the events triggering calculation of network configuration are validated. The related time period is denoted as T_L , and should be set as small as possible in the digital system in order to mimic constant validation of triggering events. Please note that T_L is different than T_M . The latter determines the history which is taken into account when validating triggering events, while the former determines the frequency at which the validation is performed (see Fig. 6.3).

With reference to Fig. 6.1, the “1. Traffic monitoring” block provides (with periodicity T_L) input data (measure of the traffic load experienced on each logical link) for the “2. Validation

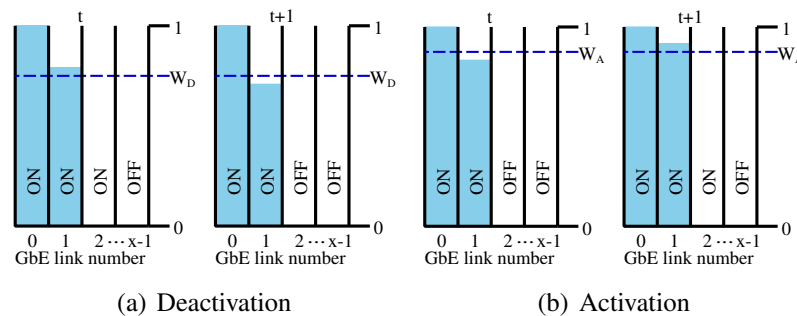


Figure 6.4: Utilization of x interfaces on a logical link operated by FUFL.

of triggering events” block. Thus, starting from the time instant t_1 , corresponding to the first validation, the n -th validation is performed at time $t_n = t_1 + (n - 1)T_L$. Each validation is performed according to the measure of the average traffic load experienced during a period T_M . Specifically, the measure provided at time t_n corresponds to the average traffic load during the period $(t_n - T_M, t_n)$. Notice that choosing a small value for T_L allows a prompt reaction to changing traffic conditions, which is particularly important during increasing traffic trend in order not to experience congestion within the network. On the other hand, the system should be as stable as possible and not follow all tiny variations of traffic. For this reason, traffic load measures are provided as averages over the period T_M , which should be chosen sufficiently long so as to hide very high frequency traffic variations.

Calculation of new network topology is triggered by violation of thresholds (W_A and W_D for FUFL, and W_L and W_H for DUFL) as explained in detail in the following subsection. This step also takes into account the stability issue by including a hysteresis cycle within the threshold mechanism.

6.5.2.3 Calculation of network configuration

We focus on two classes of approaches to calculation of network configuration introduced in Section 4.2, namely FUFL and DUFL.

FUFL: The first class is very simple and attractive for network operators [18]. It is fully distributed and involves neither changing of IP routing nor changing of the connectivity of the logical topology. The load on each GbE link constituting the logical link is monitored. We complement the FUFL algorithm from Section 4.2 with two thresholds. A GbE link is switched off when load on the previous parallel GbE link goes down below W_D (Fig. 6.4(a)). It is switched on again when the load on the previous parallel GbE link goes above W_A (Fig. 6.4(b)). W_D and W_A are defined as utilization of a GbE link. The explanation above assumes bin-packing of traffic in parallel links [123, 126]. We have not verified the load-balancing mechanisms used in the Juniper routers. If other packing (load-distribution) strategies are used, traffic on the GbE link to be switched off is shifted to other active parallel GbE link(s). According to [183, 195], link aggregation implementation in Juniper routers uses the same load-balancing algorithm as used for per-packet load balancing, i.e., the router sends successive data packets over paths without regard to individual hosts or user sessions. It uses the round-robin method to determine which path each packet takes towards its target.

DUFL: The second class DUFL is more complex, as it allows changing of IP routing, which in turn may increase the number of idle interfaces in the network, and lead even to switching off whole logical links. A logical link non-existing in the base network cannot be established though. There are many algorithms which fall into the class of DUFL (see e.g., Distributed and Adaptive Interface Switch off for Internet Energy Saving (DAISIES) [36] and Least Flow Algorithm (LFA) [61]), however their comparison is out of the scope of this chapter. The power savings in the simple base network topology from Fig. 6.2 would be basically identical.

For the sake of this work, we assume the following implementation of DUFL. The decision about an attempt to reroute traffic with the aim of deactivation or activation of a logical link is triggered by violation of the thresholds W_L and W_H , respectively. Both W_L and W_H are defined as utilization of a logical link. The traffic demands routed via E–F–G are attempted to be rerouted to link E–G if aggregated demand on the logical links E–F and F–G goes below W_L . An analogical rerouting attempt is performed when load of the logical link E–G goes below W_L .

Idle logical links with optical interfaces are switched off. The original logical topology and routing (Fig. 6.2) is restored when W_H is violated on any logical link.

6.5.2.4 Network reconfiguration


The last step concerns the application of the newly computed network configuration in network devices. To perform this step, the management system opens a telnet session on routers which need their configurations to be changed and applies the needed changes. Specifically, routing is changed and network interfaces are switched on/off according to the computed network configuration using commands specific for the Juniper routers. We ensure that rerouting is performed before a logical link is released when load decreases, and after a logical link is established when load increases.

6.5.3 Results


We parameterize the methods described in Section 6.5.2 in the following way. The thresholds are assigned with the following values: $W_D = 0.977$ and $W_A = 0.985$ for FUFL, and $W_L = 0.4885$ and $W_H = 0.9925$ for DUFL. Both T_M and T_L are set to 10 s. The chosen values are determined by the generated traffic characteristic and for the sake of demonstration of the power saving approaches. Traffic variations close to the threshold values allow us to verify the methods without waiting long (corresponding to diurnal variation of traffic).

As for the setting of T_L (the time period between two successive validations of triggering events), we had to mind the time needed for reconfiguration (Step 4 in Fig. 6.1). Specifically, T_L should be longer than the time needed for network reconfiguration (duration of Step 4). The reason is that two concurrent attempts to switch on/off a network interface could be undertaken otherwise. Due to the non-negligible values of the network reconfiguration (Step 4) experienced in our experiments, we decided to overcome this problem by not performing the validation of triggering events during network reconfiguration. In this way, it was possible to keep T_L lower than the duration of Step 4 without experiencing any concurrent attempts to switch on/off an interface.


Fig. 6.5(a) and Fig. 6.5(b) report the total power consumption and the power saving for the testbed running FUFL and DUFL on the logical topology, respectively. The total power consumption corresponds to power consumed by all active GbE interfaces together with the routers according to the data from Section 6.5.1.3. Power consumption varies more frequently with DUFL




(a) FUFL - Total power consumption (left y-axis) and power saving (right y-axis)



(b) DUFL - Total power consumption (left y-axis) and power saving (right y-axis)



(c) FUFL - Monitored traffic



(d) DUFL - Monitored traffic

Figure 6.5: FUFL and DUFL results on the testbed (mind different y-scales in (c) and (d)).

than with FUFL, since our implementation of DUFL is more aggressive in turning off the network interfaces – it attempts to switch off the whole logical links. This in turn produces in general higher power saving compared to FUFL. The difference is minor due to the simple 3-node base logical topology and IP routing schemes that we use for this demonstration (see Fig. 6.2).

Fig. 6.5(c) and Fig. 6.5(d) report the monitored traffic on logical links when FUFL and DUFL are applied, respectively. The figures report also the sine-like traffic injected to the network. As expected, all the logical links are always utilized with FUFL, and therefore each link has always aggregated traffic around 1 Gbps. On the contrary, the utilization of the links frequently changes with DUFL, since this algorithm reroutes the traffic and powers off the entire logical links. Therefore traffic on each logical link has a strong fluctuation between 0 and 2 Gbps.

The time needed for calculation of network configuration (Step 3 in Fig. 6.1) takes 0.15 s for both FUFL and DUFL. We measured also the time needed for network reconfiguration (Step 4 in Fig. 6.1). It consists of: (1) time consumed by telnet (opening a session): 11.54 s; (2) time needed to power on a GbE interface: 0.01 s; (3) time needed to power off a GbE interface: 0.01 s. We neglect the time that is needed to perform the rerouting in DUFL.

6.6 Conclusion

After looking at the challenges for EA-ARs in Chapter 5, we looked at the implementation issues for the EA-ARs in this chapter. They include not only quick calculation of the network configuration fulfilling the constraints from Section 5.1, but also methods for providing input data for the EA-ARs, and distributing the information about calculated network configuration to network devices. We discussed also the methods for traffic monitoring and implementation methods for validation of triggering events and calculation of the network configuration assuming GMPLS for network control and management. Eventually, we described how the actual implementation can be performed in an IP-over-WDM network.

We presented a survey of related work dealing with experimental activities demonstrating energy saving in backbone networks (MiDORi project and CARISMA testbed). On the top of the related work, we demonstrated the operation of energy saving approaches FUFL and DUFL on a testbed with optical GbE interfaces located in Italy (ISCOM). We showed that it was possible to dynamically adapt the network configuration to the changing load without losing traffic and with a minor increase in the packet delay. Moreover, we demonstrated that it was feasible to automatically and remotely activate and deactivate interfaces of commercial devices available today. We experienced relatively long time to reconfigure the network, which depends on how routers are accessed. We believe that this issue can be easily overcome in the operational networks using future devices designed for green networking.

Eventually, we point out that no standard and multi-vendor interface for changing the routing models of IP routers is available today. Therefore SDN is an important field of future work in the context of green networking.

7 Conclusion

Energy consumption has become a well established research topic. Approaches for energy saving are proposed for different parts of the network from the access up to the core. The focus of this work is on saving energy in core Internet Protocol (IP)-over-Wavelength Division Multiplexing (WDM) networks using Energy-Aware Adaptive Routing Solutions (EA-ARSs) utilizing sleep modes in low-demand hours.

7.1 Summary of contributions

The following contributions are included in this thesis:

- Potential of energy saving in IP-over-WDM networks in low-demand hours has been evaluated taking into account different levels of freedom of rerouting in the network (approaches Fixed Upper Fixed Lower (FUFL), Dynamic Upper Fixed Lower (DUFL) and Dynamic Upper Dynamic Lower (DUDL) described in Chapter 4). It has been shown that rerouting of IP traffic demands contributes the most to the saving of energy consumed by line cards. The additional flexibility of rerouting in the WDM layer brings barely any further savings. Furthermore, significant energy savings can be also achieved with the simple and distributed approach called FUFL, which requires only local information about the network. However, the savings provided with FUFL depend on the ratio of the maximum total demand and the capacity of a WDM channel. Also spatial distribution of demands influences the power savings achieved with FUFL. Energy consumption of a Static Base Network (SBN) over a year is estimated to be 1201 MWh (example of the Germany17 network over a year with 5 Tbps of total demand originating from measurements). Energy consumption of networks using FUFL, DUFL and DUDL (lower bound) approaches equal 785, 496, and 467 MWh, respectively.
- We defined a set of constraints that EA-ARSs need to deal with in order to be considered for implementation in operational networks (Chapter 5).
- An adaptive algorithm called Energy Watermark Algorithm (EWA) was proposed in Chapter 5 targeting switching off line cards in low-demand hours. Comparison with two other algorithms showed its good performance in terms of energy saving and reconfiguration cost. As one of few works, we considered different traffic data to design SBN and to evaluate energy savings with EWA. Therefore, we were also able to analyze the (marginal) overload experienced in the network. The sensitivity study of the EWA to its parameters showed a trade-off between the power consumption and reconfigured traffic. Smooth changes of power consumption over time and little traffic requiring rerouting are provided by the adaptive nature of EWA.

- Clear triggering events have been identified for both FUFL and EWA. The triggering events determine the moment when the EA-ARS should be launched in order to compute a new network configuration. Network configuration can start only after the computation of the new network configuration has been completed. For FUFL and EWA, the triggering events have been defined as the moments of exceeding thresholds of utilization of lightpaths.
- All the evaluation studies are performed on realistic input data. The extensive data has been made publicly available [4, 5]. The input data include extensive sets of traffic matrices with corresponding physical supply network topologies and power consumption values (Chapter 2 and Appendix D). The traffic matrices contain traffic demands between all node pairs in the network, and cover periods of days, weeks, and months with granularity of 5 min., 15 min., a day, and a month. The power data originate from product data sheets and power measurements reported in scientific papers.
- General steps needed for energy saving with EA-ARSs are identified. They include (1) traffic monitoring; (2) validation of triggering events; (3) calculation of network configuration; and (4) network reconfiguration. We discussed implementability issues faced when bringing energy saving with EA-ARSs into reality with the focus on network management and control with Generalized MultiProtocol Label Switching (GMPLS). Eventually, FUFL and DUFL have been demonstrated on an IP-over-Gigabit Ethernet (GbE) testbed (Chapter 6).
- Extensive survey of the related work has been presented distinguishing the EA-ARSs from the Energy-Aware Network Design (EA-ND) approaches. Furthermore, the layers of operation are distinguished, being either IP, or WDM, or IP and WDM.

7.2 Future work

The following topics are left for future work:

- Traffic characteristic is fundamental for energy saving with EA-ARSs. Extensive measurement campaign of traffic in commercial networks of today are desirable by the research community for several reasons. First, there are no publicly available sets of traffic matrices from a commercial network provider. Second, the traffic data from research and education networks available so far (Section 2.5) were measured in the years 2004–2005, while the services offered via the Internet have changed significantly since then (cloud services, IP telephony, Video On Demand (VOD), etc.)
- Even though many EA-ARSs have been proposed so far (see Section 3.2) and each evaluation criterion from Section 5.1 is addressed by at least one EA-ARS, there is no single EA-ARS that fulfills all the criteria. Particularly, there are only few EA-ARSs covering both IP and WDM layers. So far only one integral algorithm approach (using connectivity graph) has been identified. It does not cover protection though. As for the contribution from this thesis, EWA needs to be extended with physical layer constraints such as Routing and Wavelength Assignment (RWA), maximum length of a lightpath, and limited number of fibers on physical links. An additional constraint of no freedom of rerouting in the WDM

layer can reduce complexity of the reconfiguration with marginal influence on the energy saving achieved with EWA.

- Evaluation of EA-ARSs together with appropriate control mechanisms has not been well investigated yet.
- Demonstration of EA-ARSs addressing implementation issues challenges researchers today. Demonstrations of the proposed EA-ARSs are expected to reveal their limitations. Furthermore, successful demonstrations can encourage standardization bodies to develop standards for green networking, and product vendors to implement EA-ARSs on commercial devices. Experimental activities can be supported by Software Defined Networking (SDN). Access to an IP-over-WDM testbed is limited though.
- Development of remotely manageable network devices supporting sleep modes and having low activation and deactivation times is a prerequisite for the EA-ARSs. Their development is a separate challenge in itself.

A List of Symbols

Table A.1: Notation.

Time:	
T	set of all considered time periods (called whole period)
Δt	duration of a single time period $t \in T$
t	a single time period of duration Δt
T_{past}	set of past (known) time periods, $T_{past} \subset T$
T_{fut}	set of future (unknown) time periods, $T_{fut} \subset T$
Network:	
V	set of nodes (vertices)
E	set of physical supply links (edges)
P	set of all physical paths at which lighapths can be established
L	set of logical links
(V, E)	physical supply network
(V, L)	logical network
B	capacity of a single fiber in terms of WDM channels
W	capacity of a single WDM channel and lightpath (in bps)
$\gamma \in (0, 1]$	overprovisioning factor ratio between the lightpath's capacity used during the SBN design and its full capacity
$\delta \in (0, 1]$	maximum lightpath utilization
R	distance (in kms) between two amplifiers (span length)
N	set of IP routers that can be installed in the network
\mathcal{R}^n	capacity of router $n \in N$
W_{LCS}	number of Line Cards (LCs) that can be accommodated by a Line Card Shelf (LCS)
W_{FCS}	number of LCSs that can be accommodated by a Fabric Card Shelf (FCS)
$P_{(i,j)}$	set of all admissible physical paths in G between nodes $i \in V$ and $j \in V$
P	set of all admissible physical paths in G , that is $P = \cup_{(i,j) \in V \times V} P_{(i,j)}$

Continued on next page

Table A.1 Notation – continued from previous page.

P_i	set of all admissible paths ending at node $i \in V$, that is, $P_i = \cup_{j \in V} P(i,j)$
X_i^{LC}	number of installed line cards in each node $i \in V$ of the SBN (solution values)
Y_p^{SBN}	the set of logical links from the SBN
Q	number of line cards installed in the network (solution value)
G	number of established lightpaths in the previous time period (solution value)
Costs (Capital Expenditures (CapEx) or power):	
C^n	cost of router $n \in N$
C^e	cost of a fiber installed at physical link $e \in E$ (length dependent)
C^{LC}	cost of a lightpath (based on the cost of line cards at both ends)
C^{LCS}	cost of a Line Card Shelf (LCS)
C^{FCS}	cost of a Fabric Card Shelf (FCS)
Traffic:	
$D(t)$	traffic matrix containing all demands during time period $t \in T$
$d^{ab}(t)$	an element of a traffic matrix denoting traffic demand from node $a \in V$ and $b \in V$ during time period $t \in T$ $\sum_{(a,b) \in V \times V} d^{ab}(t)$ - total demand during period $t \in T$
D_{SBN}	traffic matrix used to design the Static Base Network (SBN)
d_{SBN}^{ab}	capacity of traffic demand between nodes $a \in V$ and $b \in V$, element of D_{SBN} ; note that different indices (namely a, b, i, j, k) indicating a node out of the set V may be used, see e.g., Eq. (4.2)
$B(t)$	traffic matrix originating from the Dwivedi-WaGner (DWG) model, containing all demands during time period $t \in T$
$b^{ab}(t)$	an element of a traffic matrix (calculated with the DWG model) denoting traffic demand from node $a \in V$ and $b \in V$ during time period $t \in T$
B_{SBN}	traffic matrix generated according to the DWG model, used to design the SBN
b^{ab}	capacity of traffic demand between nodes $a \in V$ and $b \in V$, element of B_{SBN}
K	set of commodities, corresponding to all source nodes in V
d_i^k	net demand value for commodity $k \in K$ and node $i \in V$ (SBN)
$d_i^k(t)$	net demand value for commodity $k \in K$ and node $i \in V$ during period $t \in T$
d_i^{SBN}	total demand value of network node $i \in V$ (SBN)
$d_i(t)$	total demand value of network node $i \in V$ at time $t \in T$

Continued on next page

Table A.1 Notation – continued from previous page.

$d_{ V }$	total demand value per node (SBN)
Variables:	
$f_{ij}^k, f_{ji}^k \in \mathbb{R}_+$	flow for commodity $k \in K$ between nodes $i \in V$ and $j \in V$ (in both directions, SBN)
$f_{ij}^k(t), f_{ji}^k(t) \in \mathbb{R}_+$	flow for commodity $k \in K$ between nodes $i \in V$ and $j \in V$ in time period $t \in T$ (in both directions, DUFL and DUDL)
$f_{ij}^{ab}(t) \in \{0, 1\}$	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 $j \in V$ in time period $t \in T_{fut}$ (EWA)
$f_p^{ab}(t) \in \mathbb{R}_+$	flow on path $p \in P$ for traffic demand $d^{ab}(t)$ during period $t \in T$ (FUFL)
$y_p^{SBN} \in \mathbb{Z}_+$	number of lightpaths realized on path $p \in P$ in the SBN
$y_e^{SBN} \in \mathbb{Z}_+$	number of fibers installed on physical supply link $e \in E$ (SBN)
$x_i^n \in \{0, 1\}$	decides whether to install router n at node $i \in V$ or not (SBN)
$y_p(t) \in \mathbb{Z}_+$	number of lightpaths established on the physical path $p \in P$ in time period $t \in T$
$y_l(t) \in \mathbb{Z}_+$	number of lightpaths forming logical link $l \in L$ in time period $t \in T_{fut}$ (EWA)
$x_i^{LC}(t) \in \mathbb{Z}_+$	number of line cards powered on at each node $i \in V$ in time period $t \in T_{fut}$ (EWA)
$r_{ij}^{ab}(t) \in \mathbb{R}_+$	amount of reconfigured traffic (in Gbps) between nodes a and b on the logical link from $i \in V$ to $j \in V$ in time period $t \in T_{fut}$ (GA and Eq. (5.10))
Secondary variables (postprocessing):	
$f_p^k \in \mathbb{R}_+$	flow on path $p \in P$ for $k \in K$ in the SBN, computed from f_{ij}^k, f_{ji}^k
$f_p^{ab} \in \mathbb{R}_+$	flow on path $p \in P$ for traffic demand d_{SBN}^{ab} , computed from f_p^k in the SBN
$f_p(t) \in \mathbb{R}_+$	total flow on path $p \in P$ in time period $t \in T$ (FUFL) computed from $f_p^{ab}(t)$
$R_i^{SBN} \in \mathbb{R}_+$	capacity of router installed in the node $i \in V$ in the SBN
$x_i^{LCS}(t) \in \mathbb{Z}_+$	number of LCSs used at each node $i \in V$ in time period $t \in T_{fut}$, Eq. (5.4)
$x_i^{FCS}(t) \in \mathbb{Z}_+$	number of FCSs used at each node $i \in V$ in time period $t \in T_{fut}$, Eq. (5.5)
$P^{LC}(t) \in \mathbb{R}_+$	cost (CapEx or power) of all active LCs in time period $t \in T_{fut}$, Eq. (5.2)
$P^{LCS}(t) \in \mathbb{R}_+$	cost (CapEx or power) of all active LCSs in time period $t \in T_{fut}$, Eq. (5.6)

Continued on next page

Table A.1 Notation – continued from previous page.

$P^{FCS}(t) \in \mathbb{R}_+$	cost (CapEx or power) of all active FCSs in time period $t \in T_{fut}$, Eq. (5.7)
$P^{TOT}(t) \in \mathbb{R}_+$	cost (CapEx or power) of all active LCs, LCSs, FCSs in time period $t \in T_{fut}$, Eq. (5.8)
$E^{LC} \in \mathbb{R}_+$	energy consumption due to LCs over T_{fut} , Eq. (5.3)
$E^{TOT} \in \mathbb{R}_+$	total network energy consumption (due to LCs, LCSs and FCSs) over T_{fut} , Eq. (5.9)
$\xi \in \mathbb{R}_+$	reconfiguration ratio over T_{fut} , Eq. (5.11)
$\phi \in \mathbb{R}_+$	overload ratio over T_{fut} , Eq. (5.12)
Parameters of algorithms:	
$W_D \in (0, 1]$	a threshold on the utilization of the last lightpath on a logical link to trigger deactivation of next lightpath on a logical link (FUFL)
$W_A \in [0, 1)$	a threshold on the utilization of the last lightpath on a logical link to trigger attempts to activate lightpath(s) (FUFL)
$W_L \in [0, 1)$	a threshold on the utilization of the last lightpath on a logical link to trigger attempts to release lightpath(s) (EWA and DUFL implementation on a testbed)
$W_H \in (0, 1]$	a threshold on the utilization of the last lightpath on a logical link to trigger attempts to establish additional lightpath(s) (EWA and DUFL implementation on a testbed)
$\psi \in (0, 1]$	maximum utilization of last lightpath on a logical link (EWA)
$\alpha \in [0, 1]$	weight of Genetic Algorithm (GA) fitness function, corresponding to the relative cost associated with each unit of reconfigured traffic, fitness function (5.1)
$M \in \mathbb{Z}_+$	maximum number of generations without improvements (GA)
$S \in \mathbb{Z}_+$	population size (GA)
$\kappa \in \mathbb{Z}_+$	offspring size (GA)
Traffic monitoring parameters:	
T_M	time period over which the constantly monitored traffic data is averaged and provided for evaluation triggering calculation of a new network configuration
T_L	inverse of the frequency at which the events triggering calculation of network configuration are validated

B Model variants for DUFL and DUDL

Complete mixed integer programming models to compute energy-minimal network configurations at time $t \in T$ using the two approaches Dynamic Upper Fixed Lower (DUFL) and Dynamic Upper Dynamic Lower (DUDL) are presented in this section. Both models are simplifications of the original model (4.3) used to compute the Static Base Network (SBN) in the sense that the solution space is restricted, which is achieved by fixing variables and adding constraints. Model (B.1) for DUFL and model (B.2) for DUDL are obtained from (4.3) by changing the objective, fixing variables y_e^{SBN} and x_i^n to the values of the SBN, and adding constraint (B.1e) respectively (B.2f). It follows that the installation of fibers and Internet Protocol (IP) routers is fixed to the configuration of the SBN. We are only allowed to change the IP-flow and the lightpath configuration (variables $f_{ij}^k(t)$, $f_{ij}^k(t)$, $y_p(t)$). The new objective is to minimize the number of active lightpaths. In addition the right hand side of the demand constraints (B.1b) and (B.2b) now corresponds to the demands $D(t)$ at time $t \in T$ which are by definition smaller than the maximum demands D_{SBN} used in (4.3).

$$\min \sum_{p \in P} y_p(t) \tag{B.1a}$$

$$\sum_{j \in V \setminus \{i\}} (f_{ij}^k(t) - f_{ji}^k(t)) = d_i^k(t), \quad i \in V, k \in K \tag{B.1b}$$

$$\sum_{p \in P(i,j)} W y_p(t) - \sum_{k \in K} (f_{ij}^k(t) + f_{ji}^k(t)) \geq 0, \quad (i,j) \in V \times V \tag{B.1c}$$

$$\sum_{p \in P_i} W y_p(t) \leq R_i^{SBN} - d_i(t), \quad i \in V \tag{B.1d}$$

$$y_p(t) \leq y_p^{SBN}, \quad p \in P \tag{B.1e}$$

$$f_{ij}^k(t), f_{ji}^k(t) \in \mathbb{R}_+, \quad y_p(t) \in \mathbb{Z}_+ \tag{B.1f}$$

Model (B.1) computes the energy-minimal network at time $t \in T$ using DUFL. By (B.1e) we cannot exceed the lightpath configuration in the SBN, that is, we have to use existing ones or switch them off. In particular, if a path p is not used in the SBN solution ($y_p^{SBN} = 0$) the corresponding path variable $y_p(t)$ is not generated for the DUFL model ($y_p(t) = 0$). Also notice that the fiber constraints (4.3f) become redundant because of (B.1e). The number of wavelengths per fiber can only be reduced from the value in the SBN solution. In fact there is no active physical capacity constraint for DUFL.

Model (B.2) computes the energy-minimal network at time $t \in T$ using DUDL. In contrast to model (B.1) we may establish new lightpaths as long as the number of lightpaths ending at a node is not exceeding the same number in the SBN. This guarantees that only the existing line cards are used. Notice that (B.2f) relaxes constraint (B.1e). Fiber constraints (B.2e) are not redundant for

DUDL.

$$\min \sum_{p \in P} y_p(t) \tag{B.2a}$$

$$\sum_{j \in V \setminus \{i\}} (f_{ij}^k(t) - f_{ij}^k(t)) = d_i^k(t), \quad i \in V, k \in K \tag{B.2b}$$

$$\sum_{p \in P(i,j)} W y_p(t) - \sum_{k \in K} (f_{ij}^k(t) + f_{ij}^k(t)) \geq 0, \quad (i,j) \in V \times V \tag{B.2c}$$

$$\sum_{p \in P_i} W y_p(t) \leq R_i^{SBN} - d_i(t), \quad i \in V \tag{B.2d}$$

$$\sum_{p \in P: e \in p} y_p(t) \leq B y_e^{SBN}, \quad e \in E \tag{B.2e}$$

$$\sum_{p \in P_i} y_p(t) \leq \sum_{p \in P_i} y_p^{SBN}, \quad i \in V \tag{B.2f}$$

$$f_{ij}^k(t), f_{ij}^k(t) \in \mathbb{R}_+, \quad y_p(t) \in \mathbb{Z}_+ \tag{B.2g}$$

C Details of Energy Watermark Algorithm

Three subroutines of Alg. 1 are presented in this section.

C.1 Ensuring Routability of Demands

The main requirement on an energy-saving algorithm is not to influence the connectivity of the network. This means that given a set of traffic demands, there must be at least one path in the logical topology to route each traffic demand. Therefore the first step of Energy Watermark Algorithm (EWA) is to ensure that all the demands are routable. The corresponding subroutine is shown in Alg. 4. The current network configuration, current Traffic Matrix (TM) $D(t)$, and capacity W of a lightpath are provided as input. Updated network configuration is returned by the subroutine.

In the first step, the unroutable demands are identified (line 1). Then, logical link(s) are iteratively tried to be established between the source and target nodes of the unroutable demands as long as there are unroutable demands, or it is impossible to establish corresponding logical links any more (line 3). For this reason the unroutable demands need to be sorted in a descending order according to their values (lines 4 – 5) if this has not been done so far (indicated by the demandIndex equal to 1, see lines 2 and 9). If a logical link is successfully established (line 7), then the list of unroutable demands is updated (lines 8 – 9). Note that a new logical link may change routing of several traffic demands, not only the addressed unroutable ones. If establishing of a logical link was unsuccessful, the next biggest demand is chosen (line 11) in order to establish a logical link between its source and target nodes. The algorithm terminates, when all demands are routable or it is impossible to create a logical link between source and target nodes of any unroutable demand (line 13).

Algorithm 4 Pseudo-code of ensureDemandsRoutability.

Input: netConfig, TM $D(t)$, W

Output: Updated netConfig

```
1: unrtableDemands=getUnrtableDemands(netConfig, TM);
2: demandIndex=1;
3: while (unrtableDemands.size()!=0) && (demandIndex ≤ unrtableDemands.size()) do
4:   if demandIndex==1 then
5:     unrtableDemands=sortDemands(unrtableDemands);
6:   end if
7:   if establishLL(unrtableDemands(demandIndex),  $W$ , netConfig) then
8:     unrtableDemands=getUnrtableDemands(netConfig, TM);
9:     demandIndex=1;
```

```
10:  else
11:    demandIndex++;
12:  end if
13: end while
```

C.2 Establishing Necessary Lightpaths

Once the routability of demands is ensured, the appropriate capacity of the network needs to be provided in order to meet the current traffic demands. The subroutine shown in Alg. 5 takes as input (additionally to the TM $D(t)$) current network configuration, capacity W of a lightpath, sorted list of logical links where the watermarks are exceeded, and the values of the low and high watermarks W_L and W_H . Updated network configuration and sorted list of logical links exceeding watermarks are returned as output.

The subroutine starts with the logical link with the highest utilization of the last lightpath (line 1), and tries to off-load the logical links in the decreasing order of their utilization, as long as the high watermark W_H is exceeded at the currently visited logical link LL (lines 2 – 3). All demands on the current logical link LL are identified (line 4), and a flag signaling the failure of lightpath establishment is set to true (line 5). This flag is needed to stop the attempts to establish a lightpath in the case a lightpath has already been established (see line 14) and to iterate through logical links (see line 41). Indeed the choice of the source and target of the lightpath to be established depends on the overloaded logical link LL , the TM (or precisely the value of the demand with the same source and target nodes as the LL), and the current capacity of the logical link LL (lines 6 and 12):

1. If the demand value exceeds the capacity of the logical link LL (line 6), the capacity of the LL is tried to be exceeded by establishing additional lightpath(s) in parallel to the existing ones (lines 7 – 8). In the case of a success the flag is set to false (line 9), and the sorted list of logical links exceeding watermarks is recalculated (line 10).
2. If either the demand with the same source and target as the logical link LL does not flow through LL or this demand does not exist in the TM or its value does not exceed the capacity of the logical link LL (line 12), another strategy is taken to choose lightpaths to be established to resolve violation of W_H . The demands flowing through the logical link LL are sorted from the biggest one to the smallest one (line 13). While the set of demands is not empty, and no additional lightpath has been established (line 14), the biggest (non-zero) demand flowing through the logical link LL is searched for (lines 16 – 30) under the constraint that the source and end node have sufficient number of available transmitters and receivers to establish a logical link of needed capacity (lines 19 – 20). If this requirement cannot be fulfilled, the corresponding demand is removed from the list of the demands that can potentially be offered a new (direct) lightpath (line 25). The same happens if the corresponding traffic demand does not exist in the current TM (lines 27 – 28).

Once the demand corresponding to the maximum load has been found and has a positive value (line 31), capacity of a new logical link in terms of lightpaths is calculated (line 32), and the new logical link is tried to be established (line 33). If this attempt is successful, the flag is set to false (line 34), and the sorted list of logical links exceeding watermarks

is recalculated (line 35). In any case, the corresponding demand with the same source and target as the LL attempted to be established is removed from the list of demands that can potentially be offered a new (direct) logical link (line 37).

Eventually, if no new lightpath was established in the current iteration (line 41 out of the loop in lines 2 – 46), the logical link index is decremented (line 42), that is the logical link with the next biggest utilization of the last lightpath is considered. In the case of success of the establishment of a lightpath, the logical link index is reset (line 44) to indicate again the logical link where utilization of the last lightpath is the highest.

Although the subroutine seems to be quite complex, the loops usually need to be visited only few times due to the adaptive character of the EWA. This is however highly dependent on the characteristics of the traffic, which is expected to smoothly follow the day-night pattern in the core network due to the aggregation.

Algorithm 5 Pseudo-code of establishNecessaryLpaths.

Input: netConfig, TM $D(t)$, W , sortedLLsExceedingWMs, W_L , W_H

Output: Updated netConfig and sortedLLsExceedingWMs

```

1: LLindex=sortedLLsExceedingWMs.size();
2: while (LLindex > 0) && (sortedLLsExceedingWMs(LLindex).utilLastLpath() >  $W_H$ ) do
3:   LL=sortedLLsExceedingWMs(LLindex);
4:   demands=LL.getDemandsOnLL(netConfig);
5:   failedToEstablishLpath=true;
6:   if (demands.contains(LL.id())) && (TM.contains(LL.id())) && (TM.value(LL.id()) >
   LL.capacity()) then
7:     neededCap=ceil(LL.getTotalLoad(netConfig, TM)/ $W$ ) - LL.capacity();
8:     if establishLL(LL.src(), LL.tgt(), neededCap, netConfig, TM) then
9:       failedToEstablishLpath=false;
10:      sortedLLsExceedingWMs = getSortedLLsExceedingWMs(netConfig, TM,  $W_L$ ,  $W_H$ );
11:     end if
12:   else
13:     demands=sortDemands(demands, TM);
14:     while (demands.size() > 0) && failedToEstablishLpath do
15:       maxLoad=0.0; newLLsrc=null; newLLtgt=null;
16:       for i=1; i≤demands.size(); i++ do
17:         demand=demands.get(i);
18:         if TM.contains(demand.id()) then
19:           neededCap=ceil(TM.getValue(demand.id())/ $W$ );
20:           if (TM.value(demand.id()) > maxLoad) && (demand.src().availTXs() ≥ needed-
   Cap) && (demand.tgt().availRXs() ≥ neededCap) then
21:             maxLoad=TM.value(demand.id());
22:             newLLsrc=demand.src();
23:             newLLtgt=demand.tgt();
24:           else if (demand.src().availTXs() < neededCap) || (demand.tgt().availRXs() < need-
   edCap) then
25:             demands.remove(demand);

```

```

26:         end if
27:     else
28:         demands.remove(demand);
29:     end if
30: end for
31: if maxLoad > 0 then
32:     neededCap=ceil(maxLoad/W);
33:     if establishLL(newLLsrc, newLLtgt, neededCap, netConfig, TM) then
34:         failedToEstablishLpath=false;
35:         sortedLLsExceedingWMs = getSortedLLsExceedingWMs(netConfig, TM,  $W_L$ ,
             $W_H$ );
36:     end if
37:     demands.remove(demand(newLLsrc, newLLtgt));
38: end if
39: end while
40: end if
41: if failedToEstablishLpath then
42:     LLindex--;
43: else
44:     LLindex=sortedLLsExceedingWMs.size();
45: end if
46: end while

```

C.3 Releasing Unnecessary Lightpaths

The last step of the EWA outlined in Alg. 1 is to attempt to release unnecessary lightpaths in order to save energy. Alg. 6 shows the procedure of releasing unnecessary lightpaths.

The following data is provided as input to the algorithm: current network configuration, TM $D(t)$, sorted list of logical links exceeding watermarks, the low and high watermarks (W_L and W_H) the watermarks themselves as well as the maximum utilization of the last lightpath on a logical link ψ . Updated network configuration and sorted list of logical links exceeding watermarks are returned as output.

The logical links are visited starting from the one having the lowest utilization of the last lightpath as long as there are unvisited logical links exceeding the low watermark W_L (lines 1 – 2). A flag is used in a similar way as in Alg. 5. The flag in Alg. 6 indicates the failure of the lightpath release, and not establishment. It is also initialized as true (line 3). Logical links are taken from the sorted list (line 4), and a single lightpath out of all lightpaths constituting this logical link is attempted to be released (line 5) keeping the constraint of routability of all demands in the current TM, and the constraint on the maximum utilization of the last lightpath on all the logical links ψ . If the attempt is successful, the flag is changed to true (line 6), and the sorted list of logical links exceeding watermarks is recalculated (line 7). Eventually, depending on the state of the flag (line 9), the index of the logical link is either incremented (line 10) to visit the next logical link if no lightpath has been released, or reset to 1 (line 12) to start again with the logical link with the lowest utilization of the last lightpath if a lightpath has been released.

Algorithm 6 Pseudo-code of releaseUnnecessaryLpaths.

Input: netConfig, TM $D(t)$, sortedLLsExceedingWMs, W_L , W_H , ψ **Output:** Updated netConfig and sortedLLsExceedingWMs

```
1: LLindex=1;
2: while (LLindex  $\leq$  sortedLLsExceedingWMs.size()) &&
   (sortedLLsExceedingWMs(LLindex).utilLastLpath()  $<$   $W_L$ ) &&
   (sortedLLsExceedingWMs.size()  $>$  0) do
3:   failedToReleaseLpath=true;
4:   LL=sortedLLsExceedingWMs(LLindex);
5:   if releaseLastLpath(LL, netConfig, TM,  $\psi$ ) then
6:     failedToReleaseLpath=false;
7:     sortedLLsExceedingWMs = getSortedLLsExceedingWMs(netConfig, TM,  $W_L$ ,  $W_H$ );
8:   end if
9:   if failedToReleaseLpath then
10:    LLindex++;
11:  else
12:    LLindex=1;
13:  end if
14: end while
```

D Power consumption of single network elements

We present an overview of raw data available in the product data sheets. The content of this appendix is an extension of [5]. Some further data covering also the Optical Transport Network (OTN) and Ethernet layers are available in [6, 11], however they are out of scope of this work.

D.1 IP devices

Table D.1: Power consumption of line cards in the IP layer.

Device	Router	Power [Watts]	Ref.
4 port GE	Cisco GSR 12008s	92 (no traffic), 106 (maximum)	[19, 196]
1 port OC-12/POS	Cisco GSR 12008s	72 (no traffic)	[19]
1 port OC-48/POS	Cisco GSR 12008s	70 (no traffic)	[19]
1 port Fast Ethernet	Cisco 7507	26 (no traffic)	[19]
1 port Gigabit Ethernet	Cisco 7507	30 (no traffic)	[19]
1 port 1.544 Mb/s DS1	Cisco 7507	49 (no traffic)	[19]
10-Gigabit Ethernet Physical Layer Interface Module with eight optics modules + Modular Services Card	Cisco CRS-1	150 + 350 = 500	[30, 31]
10-Gigabit Ethernet WDM-PHY Physical Layer Interface Module + Modular Services Card	Cisco CRS-1	150 + 350 = 500	[30, 31]
OC-768c/STM-256c POS Physical Layer Interface Module + Modular Services Card	Cisco CRS-1	75 + 350 = 425	[31, 197]
OC-768c/STM-256c WDM-POS Physical Layer Interface Module + Modular Services Card	Cisco CRS-1	150 + 350 = 500	[31, 197]

Continued on next page

Table D.1 Power consumption of line cards in the IP layer – continued from previous page.

Device	Router	Power [Watts]	Ref.
OC-192c/STM-64c POS/DPT Physical Layer Interface Module + Modular Services Card	Cisco CRS-1	138 + 350 = 488	[31, 197]
OC-48c/STM-16c POS/DPT Physical Layer Interface Module + Modular Services Card	Cisco CRS-1	150 + 350 = 488 (maximum) or 136 + 350 = 486 (derated)	[31, 197]
unspecified	Cisco CRS-1	approx. 500	[198]
unspecified	Cisco CRS-1 (8- slot)	1000 on average	[34]
10 Gbps	unspecified	200	[199]
unspecified	unspecified	100	[28]
10-Gigabit Ethernet DWDM OTN PIC	Juniper T1600, T640, T320	26.6	[200– 202]
10-Gigabit Ethernet DWDM PIC	Juniper T1600, T640, T320	26.6	[200– 202]
10-Gigabit Ethernet LAN/WAN PIC with XFP	Juniper T1600, T640	43	[200, 201]
10-Gigabit Ethernet PIC with XENPAK	Juniper T1600, T640, T320	26.6	[200– 202]
10-Gigabit Ethernet En- hanced IQ2 (IQ2E) PIC with XFP	Juniper T1600, T640, T320	56	[200– 202]
10-Gigabit Ethernet IQ2 PIC with XFP	Juniper T1600, T640, T320	56	[200– 202]
SONET/SDH OC768c/STM256 PIC	Juniper T1600, T640	65.7	[200, 201]
SONET/SDH OC192/STM64 PICs with XFP	Juniper T1600, T640	25 (one-port), 53.1 (four- port)	[200, 201]
SONET/SDH OC192/STM64 PIC with XFP	Juniper T320	25 (one-port)	[202]
SONET/SDH OC192c/STM64 PIC	Juniper T1600, T640, T320	21.6	[200– 202]

Table D.2: Power consumption of shelves in the IP layer.

Device	Router	Power [Watts]	Ref.
5 configurable interface slots, 3 rack chassis	Cisco 7507	max. 945, typical 650	[203]
8 slots	Cisco GSR 12008	1620	[204]
16-Slot Single-Shelf System	Cisco CRS-1	max. 10920 fully configured	[7]
8-Slot Single-Shelf System	Cisco CRS-1	max. 5992 fully configured	[7]
4-Slot Single-Shelf System	Cisco CRS-1	4326@16200 BTU/HR ¹	[7]
Fabric Card Chassis	Cisco CRS-1	max. 9100@31050 BTU/HR	[7]
72 line card shelves interconnected by 8 fabric shelves	Cisco CRS-1	1020000	[7, 205, 206]
single-rack (640 Gbps)	Cisco CRS-1	10900	[207]
Juniper T320 (320 Gbps)	Juniper T Series	2880 (agency label, max), 2314 (theoretical aggregate)	[8]
Juniper T640 (640 Gbps)	Juniper T Series	7296 (agency label, max), 4517 (theoretical aggregate)	[8]
Juniper T1600 (1.6 Tbps)	Juniper T Series	8352 (agency label, max), 7008 (theoretical aggregate)	[8]
Juniper TX Matrix with 4 x T640 (2.5 Tbps)	Juniper T Series	4560 (agency label, max), 2976 (theoretical aggregate)	[8]
Juniper TX Matrix Plus with 16 x T1600 (25.6 Tbps)	Juniper T Series	12750 (preliminary)	[8]

According to [7, 205, 206], each complete CRS-1 system consumes 1020 kW, which shows a significant potential of energy savings.

The power consumption of CRS-1 chassis (single- and multi-shelves of different capacities in terms of 40 Gbps line cards) is presented in Table D.3. According to [7], the power consumption of a 16-Slot Single-Shelf-System (built from the single line card chassis (shelf)) fully configured with line cards with traffic running is equal to 10.92 kW. Since the power of each line card is 500 W [30, 31, 198], we assume that the power of an idle 16-slot Single-Shelf-System without line cards is equal to 2.92ZkW.

In order to increase the capacity of the CRS-1 Carrier Routing system, the above mentioned line card shelves can be interconnected into a Multishelf System using fabric card shelves. Each fabric card shelf can accommodate up to 9 line card shelves. Altogether a Multishelf System can contain up to 72 Line Card Shelf (LCS) interconnected by eight Fabric Card Shelves (FCSs) (we skip some configurations of the Multishelf System in Table D.3). Power of a FCS (chassis) is equal to 9.1 kW @ 31.050 BTU/hr [7].

¹BTU = BritishThermalUnit = 0.2931Wh; BTU/HR \simeq 0.293W

Table D.3: Power consumption of chassis in multi-shelf systems of Cisco CRS-1 Carrier Routing System using 16-Slot Single-Shelf Systems [7].

Shelftype	Number of 40 Gbps line card slots	Number of fabric card shelves needed	Power [kW]
SH-IP-640	16	0	$1 \cdot 2.92 + 0 \cdot 9.1 = 2.92$
SH-IP-1280	32	1	$2 \cdot 2.92 + 1 \cdot 9.1 = 14.94$
SH-IP-1920	48	1	$3 \cdot 2.92 + 1 \cdot 9.1 = 17.86$
SH-IP-2560	64	1	$4 \cdot 2.92 + 1 \cdot 9.1 = 20.78$
SH-IP-3200	80	1	$5 \cdot 2.92 + 1 \cdot 9.1 = 23.7$
SH-IP-3840	96	1	$6 \cdot 2.92 + 1 \cdot 9.1 = 26.62$
SH-IP-4480	112	1	$7 \cdot 2.92 + 1 \cdot 9.1 = 29.54$
SH-IP-5120	128	1	$8 \cdot 2.92 + 1 \cdot 9.1 = 32.46$
SH-IP-5760	144	1	$9 \cdot 2.92 + 1 \cdot 9.1 = 35.38$
SH-IP-6400	160	2	$10 \cdot 2.92 + 2 \cdot 9.1 = 47.40$
SH-IP-7040	176	2	$11 \cdot 2.92 + 2 \cdot 9.1 = 50.32$
...
SH-IP-11520	288	2	$18 \cdot 2.92 + 2 \cdot 9.1 = 70.76$
SH-IP-12160	304	3	$19 \cdot 2.92 + 3 \cdot 9.1 = 82.78$
...
SH-IP-17280	432	3	$27 \cdot 2.92 + 3 \cdot 9.1 = 106.14$
SH-IP-17920	448	4	$28 \cdot 2.92 + 4 \cdot 9.1 = 118.16$
...
SH-IP-23040	576	4	$36 \cdot 2.92 + 4 \cdot 9.1 = 141.52$
SH-IP-23680	592	5	$37 \cdot 2.92 + 5 \cdot 9.1 = 118.16$
...
SH-IP-28800	720	5	$45 \cdot 2.92 + 5 \cdot 9.1 = 176.90$
SH-IP-29440	736	6	$46 \cdot 2.92 + 6 \cdot 9.1 = 188.92$
...
SH-IP-34560	864	6	$54 \cdot 2.92 + 6 \cdot 9.1 = 212.28$
SH-IP-35200	880	7	$55 \cdot 2.92 + 7 \cdot 9.1 = 224.30$
...
SH-IP-40320	1008	7	$63 \cdot 2.92 + 7 \cdot 9.1 = 247.66$
SH-IP-40960	1024	8	$64 \cdot 2.92 + 8 \cdot 9.1 = 259.68$
...
SH-IP-46080	1152	8	$72 \cdot 2.92 + 8 \cdot 9.1 = 283.04$

As for the Juniper T Series Core routers, apart from the Physical Interfaces Cards (PICs) mentioned in Table D.1, the power consumption of various other PICs can be found in [200–202]. It needs to be mentioned that the PICs are Short Reach (SR) [200], which explains that their power consumption is lower by an order of magnitude compared to their CISCO counterparts. Moreover, Juniper does not mention an equivalent of CISCO’s Modular Services Card . However, in the appendices of the hardware guides [208–211] they give detailed current values to estimate power

consumption of certain configuration of Juniper routers. We assume for both CISCO and Juniper that the power values indicated in the data sheets [7, 8] include the power consumed by line cards, although it is not clearly explained there (especially in [8]).

Table D.4 shows power consumption of Juniper T1600 core router(s) interconnected by a TX Matrix Plus. The TX Matrix Plus can interconnect up to 16 T1600 chassis, what corresponds to 256 40 Gbps ports [8]. Power consumption of a single Juniper T1600 core router equals 7008 W. We calculate the power of a T1600 chassis by subtracting power of 16 OC768c/STM256 PICs (65.7 W [200]) from this value, which results in 5956.8 W. For completeness, we repeat from Table D.2 that the power consumption of a TX Matrix Plus is 12750ZW.

Table D.4: Power consumption of chassis in multi-shelf systems of Juniper T1600 core routers using Juniper TX Matrix [8].

Shelftype	Number of 40 Gbps line card slots	Number of fabric card shelves needed	Power [kW]
SH-IP-640	16	0	$1 \cdot 5.9568 + 0 \cdot 12.75 = 5.9568$
SH-IP-1280	32	1	$2 \cdot 5.9568 + 1 \cdot 12.75 = 24.6636$
SH-IP-1920	48	1	$3 \cdot 5.9568 + 1 \cdot 12.75 = 30.6204$
SH-IP-2560	64	1	$4 \cdot 5.9568 + 1 \cdot 12.75 = 36.5772$
SH-IP-3200	80	1	$5 \cdot 5.9568 + 1 \cdot 12.75 = 42.534$
SH-IP-3840	96	1	$6 \cdot 5.9568 + 1 \cdot 12.75 = 48.4908$
SH-IP-4480	112	1	$7 \cdot 5.9568 + 1 \cdot 12.75 = 54.4476$
SH-IP-5120	128	1	$8 \cdot 5.9568 + 1 \cdot 12.75 = 60.4044$
SH-IP-5760	144	1	$9 \cdot 5.9568 + 1 \cdot 12.75 = 66.3612$
SH-IP-6400	160	1	$10 \cdot 5.9568 + 1 \cdot 12.75 = 72.318$
SH-IP-7040	176	1	$11 \cdot 5.9568 + 1 \cdot 12.75 = 78.2748$
SH-IP-7680	192	1	$12 \cdot 5.9568 + 1 \cdot 12.75 = 84.2316$
SH-IP-8320	208	1	$13 \cdot 5.9568 + 1 \cdot 12.75 = 90.1884$
SH-IP-8960	224	1	$14 \cdot 5.9568 + 1 \cdot 12.75 = 96.1452$
SH-IP-9600	240	1	$15 \cdot 5.9568 + 1 \cdot 12.75 = 102.102$
SH-IP-10240	256	1	$16 \cdot 5.9568 + 1 \cdot 12.75 = 108.0588$

D.2 WDM devices

Table D.5: Power consumption of transponders / muxponders in WDM layer.

Device	Power [Watts]	Ref.
2.5-GBPS Multirate Transponder Card for the CISCO ONS 15454 Multiservice Transport Platform	Typical 25, maximum 35	[212]
Continued on next page		

Table D.5 – continued from previous page.

Device	Power [Watts]	Ref.
10-GBPS Multirate Transponder Card for the CISCO ONS 15454 Multiservice Transport Platform	Typical 45, maximum 50	[213]
4 x 2.5-GBPS Muxponder Card for the CISCO ONS 15454 Multiservice Transport Platform	Typical 45, maximum 50	[214]
10-Gbps Full-Band Tunable Multirate Enhanced Transponder Card for the Cisco ONS 15454 Multiservice Transport Platform	Typical 35, maximum 50	[215]
Cisco 2-Channel SFP WDM Transponder	maximum 9.7	[216]
Lucent Metropolis Multi-Rate Transponder for Service Providers	7RU shelf: 50 per fan, 330 max (fully populated shelf)	[217]
Ciena F10-T 10G Transponder Module For the CN 4200 FlexSelect Advanced Services Platform Family	35 (with maximum FEC), 41 (tunable, with maximum FEC)	[218]
Fujitsu Flashwave 7200 - Transponder, Protection, Regenerator System	ANSI shelf: 381 typical for 16 OC-48/STM-16 transponders, 333 typical for 8 OC-192/STM-64 transponders; ETSI shelf: 334 typical for 11 OC-48/STM-16 transponders, 292 typical for 7 OC-192/STM-64 transponders	[219]
Fujitsu Flashwave 7300 - Transponder, Protection, Regenerator System	681 for 18 bidirectional OC-192s/STM-64s	[220]
Alcatel-Lucent Wave Star OLS 1.6T ultra-long-haul systems	73	[34]
Fujitsu Flashwave 7500 small system (compact ROADM for Mid-Size Applications)	<440@1500 BTU/HR (OLC shelf (full))	[221]

Table D.6: Power consumption of regenerators in WDM layer.

Device	Power [Watts]	Ref.
Cisco Optical Regenerator (OC-48/STM-16 Bidirectional)	max. 100@342 BTU/HR	[222]
JDSU WaveShifter 200 ER 3R	typical 4.5, max. 6	[223]
iLynx STM-16 / OC-48 Regenerator	nominal 8	[224]

Table D.7: Power consumption of WDM terminals in WDM layer.

Device	Power [Watts]	Ref.
connecting edge nodes to the core node	1500 W for every 64 wavelengths	[205, 206]
connecting core nodes (Fujitsu Flashwave 7700)	811 W for every 176 channels	[205, 206, 225, 226]
undersea terminal system	9000 W 4 fiber pairs with 64 wavelengths at 10 Gbps	[207]

Table D.8: Power consumption of OLAs in WDM layer.

Device	Power [Watts]	Ref.
multiple wavelength amplifier (1 amplifier per ca. 200 km)	6 W per fiber	[205, 206]
CISCO ONS 15501 EDFA	15 maximum, 8 typical	[34, 227]
intermediate line amplifier (1 ILA per ca. 108 km), Fujitsu Flashwave 7700	622 W for every 176 channels	[205, 206, 225, 226]
Fujitsu Flashwave 7600 (in-line amplifier)	601 W typical for 32 wavelengths	[228]
Lucent Metropolis Multi-wavelength Optical Amplifier (1 OA per 140 km in single span, 1 OA per 100 km in multiple spans)	typical 6, max. 25	[229]
JDSU OA 400 Amplifier Series - Compact, Low Cost Single-Channel or Narrow Band Amplifiers	4.5	[230]
infinera Optical Line Amplifier	205 (typical)	[231]
single amplifier (every 70 km)	1000	[28]
undersea repeater	40 (every 50 km)	[207]

Table D.9: Power consumption of optical switches in WDM layer.

Device	Power [Watts]	Ref.
Ciena Configurable OADM For Core-Stream Agility Optical Transport System	typical 250	[232]
Fujitsu Flashwave 7500 Metro/Regional ROADM Platform	<360@1228 BTU/HR (Optical/HUB shelf); <680@2318 BTU/HR (2D-ROADM shelf); <930@3170 BTU/HR (Tributary shelf)	[233]
Fujitsu Flashwave 7500 small system (compact ROADM for Mid-Size Applications)	<240@820 BTU/HR (OADM shelf (typical))	[221]
Fujitsu Flashwave 7600 (Regional Long Haul DWDM Solution)	443 typical for 32 wavelengths	[228]
Finisar Dynamic Wavelength Processor DWP 100 Wavelength Selective Switch	typical 20	[234]
Finisar Dynamic Wavelength Processor DWP 50 Wavelength Selective Switch	typical 20	[235]
Remote Optical Add/Drop Multiplexer/Photonic Cross Connect (Implemented current equipment)	40 / 100 Gbps at 10 Gbps in 50 GHz spacing	[236]
Remote Optical Add/Drop Multiplexer/Photonic Cross Connect (Next generation equipment)	10 / 100 Gbps at 40 Gbps in 50 GHz spacing	[236]
Infinera DTN Switched WDM System (ROADM)	1500 (typical, fully loaded), 2730 (max)	[237]
128 x 128 Calient OXC	<100	[207]
Calient DiamondWave FiberConnect - Fiber Optic Cross-Connection System	0.47 per connection (fully loaded)	[238]
Calient DiamondWave PXC - photonic switching system for metro and long-haul network environment (128x128 or 256x256)	0.47 per connection (fully loaded)	[239]
MRV OCC - Optical Cross Connect - Single-Mode Fiber Optic Switch	<750 (fully loaded)	[240]
MRV OCC 96 - Optical Cross Connect	55 (max)	[241]

Significant differences of power consumption of comparable network elements can be observed. E.g., an OC-48/STM-16 regenerator by Cisco [222] consumes max. 100 W, while nominal power consumption of the iLynx OC-48/STM-16 is 8 W [224]. Another example is the power of a transponder, which ranges from 9.7 W for Cisco 2-Channel SFP WDM Transponder [216] to 73 W for Alcatel-Lucent Wave Star OLS 1.6T ultra-long-haul systems [34]. These differences may be caused by the differences in technologies used for deployment of the network elements, as well as

the technological improvement of energy-efficiency over time.

D.3 Other aspects regarding power consumption

Network elements not only consume energy, but also generate heat. Following [205], we account one Watt spent on cooling for every Watt spent by any of the network devices.

Another issue is the dependency of the energy consumption on the network load. According to [19], traffic has an impact on the power consumption, however the costs of powering the chassis and the line cards dominate the overall power profile of a network. In the scenario considered in [19], traffic (of various characteristics) increased the power consumption of Cisco GSR 12000 with a 4-port Gigabit Ethernet engine 3 line card and a 1-port OC-48/POS line card (the line cards contain transceivers [196, 204, 242]) by about 20 Watts (out of the 756 Watts consumed in the idle case), which corresponds to only 2.65% of the power consumed by the router.

E Networks – nodes and physical supply links

The Tables E.1 – E.5 contain details about nodes and links of the physical supply networks considered in this work and available in [4]. The ids of demand node pairs are created in the following way: $sourceId \cdot numberOfNodes + targetId$.

Table E.1: Network nodes with corresponding locations and node ids for the Géant and Germany17 networks.

(a) Géant network			(b) Germany17 network		
Node name	Corresponding country	Node Id	Node name	Corresponding city	Node Id
at1.at	Austria	0	Berlin	Berlin	0
be1.be	Belgium	1	Bremen	Bremen	1
ch1.ch	Switzerland	2	Dortmund	Dortmund	2
cz1.cz	Czech Republic	3	Duesseldorf	Duesseldorf	3
de1.de	Germany	4	Essen	Essen	4
es1.es	Spain	5	Frankfurt	Frankfurt	5
fr1.fr	France	6	Hamburg	Hamburg	6
gr1.gr	Greece	7	Hannover	Hannover	7
hr1.hr	Croatia	8	Karlsruhe	Karlsruhe	8
hu1.hu	Hungary	9	Koeln	Koeln	9
ie1.ie	Ireland	10	Leipzig	Leipzig	10
il1.il	Israel	11	Mannheim	Mannheim	11
it1.it	Italy	12	Muenchen	Muenchen	12
lu1.lu	Luxembourg	13	Norden	Norden	13
nl1.nl	Netherlands	14	Nuernberg	Nuernberg	14
ny1.ny	New York City	15	Stuttgart	Stuttgart	15
pl1.pl	Poland	16	Ulm	Ulm	16
pt1.pt	Portugal	17			
se1.se	Sweden	18			
si1.si	Slovenia	19			
sk1.sk	Slovakia	20			
uk1.uk	United Kingdom	21			

Table E.2: Network nodes with corresponding locations and node ids for the Abilene network.

Node name	Corresponding city	Node Id
ATLAM5	Atlanta M5	0
ATLAng	Atlanta	1
CHINg	Chicago	2
DNVRng	Denver	3
HSTNng	Houston	4
IPLSng	Indianapolis	5
KSCYng	Kansas City	6
LOSAng	Los Angeles	7
NYCMng	New York City	8
SNVAng	Sunnyvale	9
STTLng	Seattle	10
WASHng	Washington	11

Table E.3: Physical supply links in the Abilene network.

Link name	Source id	Target id	Length [km]	Link id
ATLAM5_ATLAng	0	1	2	0
ATLAng_HSTNng	1	4	1143	1
ATLAng_IPLSng	1	5	690	2
ATLAng_WASHng	1	11	874	3
CHINg_IPLSng	2	5	260	4
CHINg_NYCMng	2	8	1147	5
DNVRng_KSCYng	3	6	745	6
DNVRng_SNVAng	3	9	1516	7
DNVRng_STTLng	3	10	1573	8
HSTNng_KSCYng	4	6	1028	9
HSTNng_LOSAng	4	7	2196	10
IPLSng_KSCYng	5	6	903	11
LOSAng_SNVAng	7	9	505	12
NYCMng_WASHng	8	11	336	13
SNVAng_STTLng	9	10	1138	14

Table E.4: Physical supply links in the Géant network.

Link name	Source id	Target id	Length [km]	Link id
at1.at_ch1.ch	0	2	805	0
at1.at_de1.de	0	4	599	1
at1.at_hu1.hu	0	9	218	2
at1.at_ny1.ny	0	15	6803	3
at1.at_si1.si	0	19	278	4
be1.be_fr1.fr	1	6	264	5
be1.be_lu1.lu	1	13	188	6
be1.be_nl1.nl	1	14	169	7
ch1.ch_fr1.fr	2	6	411	8
ch1.ch_it1.it	2	12	251	9
cz1.cz_de1.de	3	4	412	10
cz1.cz_pl1.pl	3	16	310	11
cz1.cz_sk1.sk	3	20	291	12
de1.de_fr1.fr	4	6	479	13
de1.de_gr1.gr	4	7	1795	14
de1.de_ie1.ie	4	10	1089	15
de1.de_it1.it	4	12	520	16
de1.de_nl1.nl	4	14	359	17
de1.de_se1.se	4	18	1185	18
es1.es_fr1.fr	5	6	1054	19
es1.es_it1.it	5	12	1189	20
es1.es_pt1.pt	5	17	505	21
fr1.fr_lu1.lu	6	13	288	22
fr1.fr_uk1.uk	6	21	344	23
gr1.gr_it1.it	7	12	1454	24
hr1.hr_hu1.hu	8	9	305	25
hr1.hr_si1.si	8	19	116	26
hu1.hu_sk1.sk	9	20	164	27
ie1.ie_uk1.uk	10	21	464	28
il1.il_it1.it	11	12	2659	29
il1.il_nl1.nl	11	14	3297	30
nl1.nl_uk1.uk	14	21	360	31
ny1.ny_uk1.uk	15	21	5575	32
pl1.pl_se1.se	16	18	778	33
pt1.pt_uk1.uk	17	21	1589	34
se1.se_uk1.uk	18	21	1427	35

Table E.5: Physical supply links in the Germany17 network.

Link name	Source id	Target id	Length [km]	Link id
Berlin_Hamburg	0	6	255	0
Berlin_Hannover	0	7	251	1
Berlin_Leipzig	0	10	152	2
Bremen_Hamburg	1	6	100	3
Bremen_Hannover	1	7	103	4
Bremen_Norden	1	13	121	5
Dortmund_Essen	2	4	35	6
Dortmund_Hannover	2	7	187	7
Dortmund_Koeln	2	9	74	8
Dortmund_Norden	2	13	234	9
Duesseldorf_Essen	3	4	29	10
Duesseldorf_Koeln	3	9	38	11
Frankfurt_Hannover	5	7	263	12
Frankfurt_Koeln	5	9	146	13
Frankfurt_Leipzig	5	10	295	14
Frankfurt_Mannheim	5	11	74	15
Frankfurt_Nuernberg	5	14	191	16
Hamburg_Hannover	6	7	131	17
Hannover_Leipzig	7	10	213	18
Karlsruhe_Mannheim	8	11	54	19
Karlsruhe_Stuttgart	8	15	61	20
Leipzig_Nuernberg	10	14	230	21
Muenchen_Nuernberg	12	14	149	22
Muenchen_Ulm	12	16	119	23
Nuernberg_Stuttgart	14	15	164	24
Stuttgart_Ulm	15	16	74	25

F Publications

Publications that contain parts of the results presented in this thesis are listed below. The publications containing major contributions to this work are marked with a *.

Journal Articles:

- [136]* F. Idzikowski, S. Orłowski, C. Raack, H. Woesner, and A. Wolisz. Dynamic routing at different layers in IP-over-WDM networks – maximizing energy savings. *Optical Switching and Networking, Special Issue on Green Communications*, 8(3):181–200, July 2011.
 - [6] 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(2):86–102, October 2012.
 - [137]* 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.
 - [137] A. Ahmad, A. Bianco, E. Bonetto, L. Chiaraviglio, and F. Idzikowski. Energy-Aware Design of Multi-layer Core Networks. *Journal of Optical Communications and Networking*, 5(10):A127–A143, October 2013.
 - [112]* 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. *IEEE Communications Magazine*, 51(11):94–104, November 2013.
 - [182]* 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, Special Issue on Innovative Techniques for Power Consumption Saving in Telecommunication Networks*, 2014(806960):1–11, March 2014.
-

Conference Articles:

- [152]* F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz. Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios. In Proc. of the ONDM, Kyoto, Japan, February 2010.
 - [243] E. Bonetto, L. Chiaraviglio, D. Cuda, F. Idzikowski, and F. Neri. Exploiting Traffic Dynamics in Power-Aware Logical Topology Design. In Proc. of the ECOC, Geneva, Switzerland, September 2011.
 - [140] F. Idzikowski, L. Chiaraviglio, and F. Portoso. Optimal Design of Green Multi-layer Core Networks. In Proc. of the e-Energy, Madrid, Spain, May 2012.
 - [176]* F. Idzikowski, L. Chiaraviglio, and E. Bonetto. EWA: an Adaptive Algorithm for Energy Saving in IP-over-WDM Networks. In Proc. of the NOC, Vilanova i la Geltru, Spain, June 2012.
 - [244] F. Idzikowski, R. Duque, E. Le Rouzic, F. Jiménez, L. Chiaraviglio, and M. Ajmone Marsan. Energy saving in optical operator networks: the challenges, the trend vision, and some results. In Proc. of the ECOC, Amsterdam, The Netherlands, September 2012.
 - [20] W. Van Heddeghem, F. Idzikowski, E. Le Rouzic, J. Y. Mazeas, H. Poignant, S. Salaun, B. Lannoo, and D. Colle. Evaluation of Power Rating of Core Network Equipment in Practical Deployments. In Proc. of the greencom, online conference, September 2012.
 - [18]* 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 Proc. of the ICC, Budapest, Hungary, June 2013.
-

Technical Reports:

- [5]* F. Idzikowski. Power consumption of network elements in IP over WDM networks. Technical Report TKN-09-006, Technical University of Berlin, Telecommunication Networks Group, July 2009.
 - [11] W. Van Heddeghem and F. Idzikowski. Equipment power consumption in optical multilayer networks – source data. Technical Report IBCN-12-001-01, Ghent University, Department of Information Technology (INTEC), IBBT, January 2012.
 - [180]* F. Idzikowski, E. Bonetto, and L. Chiaraviglio. EWA – an adaptive algorithm using watermarks for energy saving in IP-over-WDM networks. Technical Report TKN-12-002, Technical University of Berlin, Telecommunication Networks Group, May 2012.
-

G Bibliography

- [1] M. Herzog, “RINGOSTAR – an evolutionary performance upgrade of optical ring networks,” PhD thesis, Technische Universität Berlin, April 2006.
- [2] 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*, vol. 7, no. 9, pp. 814–833, September 2008.
- [3] A. Morea, J. Perelló, S. Spadaro, D. Verchère, and M. Vigoureux, “Protocol enhancements for ‘greening’ optical networks,” *Bell Labs Technical Journal*, vol. 18, no. 3, pp. 211–230, December 2013.
- [4] “Sndlib: library of test instances for survivable fixed telecommunication network design,” <http://sndlib.zib.de/home.action>, October 2012.
- [5] F. Idzikowski, “Power consumption of network elements in IP over WDM networks,” Technical University of Berlin, Telecommunication Networks Group, Tech. Rep. TKN-09-006, July 2009.
- [6] W. Van Heddeghem, F. Idzikowski, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, “Power consumption modeling in optical multilayer networks,” *Photonic Network Communications*, vol. 24, no. 2, pp. 86–102, October 2012.
- [7] Cisco, “Cisco CRS-1 Production Brochure,” http://www.cisco.com/en/US/prod/collateral/routers/ps5763/prod_brochure0900aecd800f8118.pdf, not publicly available any more, October 2008.
- [8] Product Data Sheet, “Juniper T Series Core Routers: T320, T640, T1600, TX Matrix, and TX MATRIX PLUS,” <http://www.juniper.net/us/en/local/pdf/datasheets/1000051-en.pdf>, March 2009.
- [9] 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 *Proc. of 2nd International Symposium on Advanced Networks and Telecommunication Systems (ANTS), Mumbai, India*, December 2008, pp. 1–3.
- [10] 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*, vol. 17, no. 2, pp. 285–295, March–April 2011.
- [11] W. Van Heddeghem and F. Idzikowski, “Equipment power consumption in optical multilayer networks – source data,” Ghent University, Department of Information Technology (INTEC), IBBT, Tech. Rep. IBCN-12-001-01, January 2012.

- [12] M. C. Parker and S. D. Walker, "Roadmapping ICT: An absolute energy efficiency metric," *Journal of Optical Communications and Networking*, vol. 3, no. 8, pp. A49–A58, August 2011.
- [13] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, "Energy Efficiency in Telecom Optical Networks," *IEEE Communications Surveys & Tutorials*, vol. 12, no. 4, pp. 441–458, 2010.
- [14] 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," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 2, pp. 223–244, 2011.
- [15] A. P. Bianzino, C. Chaudet, D. Rossi, and J.-L. Rougier, "A Survey of Green Networking Research," *IEEE Communications Surveys & Tutorials*, vol. 14, no. 1, pp. 3–20, 2012.
- [16] F. Musumeci, M. Tornatore, and A. Pattavina, "A power consumption analysis for IP-over-WDM core network architectures," *Journal of Optical Communications and Networking*, vol. 4, no. 2, pp. 108–117, February 2012.
- [17] R. Bolla, F. Davoli, R. Bruschi, K. Christensen, F. Cucchietti, and S. Singh, "The potential impact of green technologies in next-generation wireline networks: Is there room for energy saving optimization?" *IEEE Communications Magazine*, vol. 49, no. 8, pp. 80–86, August 2011.
- [18] 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 *Proc. of the ICC, Budapest, Hungary*, June 2013.
- [19] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," in *Proc. of the INFOCOM, Phoenix, USA*, April 2008.
- [20] W. Van Heddeghem, F. Idzikowski, E. Le Rouzic, J. Y. Mazeas, H. Poignant, S. Salaun, B. Lannoo, and D. Colle, "Evaluation of Power Rating of Core Network Equipment in Practical Deployments," in *Proc. of the greencom, online conference*, September 2012.
- [21] A. Morea, S. Spadaro, O. Rival, J. Perelló, F. Agraz, and D. Verchère, "Power management of optoelectronic interfaces for dynamic optical networks," in *Proc. of the ECOC, Geneva, Switzerland*, September 2011.
- [22] B. Mukherjee, *Optical WDM networks*. Springer, 2006.
- [23] W. Stallings, *Data and Computer Communications*. Pearson Prentice Hall, 2007.
- [24] A. Farrel and I. Bryskin, *GMPLS Architecture and Applications*. Morgan Kaufmann, 2006.
- [25] A. Morea, J. Perelló, F. Agraz, and S. Spadaro, "Demonstration of GMPLS-controlled device power management for next generation green optical networks," in *Proc. of the OFC, Los Angeles, USA*, March 2012.

- [26] E. Bonetto, M. Mellia, and M. Meo, “Energy profiling of ISP points of presence,” in *Proc. of the ICC, Ottawa, Canada*, June 2012.
- [27] A. Alimian, B. Nordman, and D. Kharitonov, “Network and telecom equipment – energy and performance assessment,” http://www.ecrinitiative.org/pdfs/ECR_3_0_1.pdf, December 2010.
- [28] L. Chiaraviglio, M. Mellia, and F. Neri, “Energy-aware backbone networks: a case study,” in *Proc of the ICC workshop on Green Communications, Dresden, Germany*, June 2009.
- [29] “Powerlib: ICT networks - power consumption reference values and database,” <http://powerlib.intec.ugent.be/database/>, April 2014.
- [30] Product Installation Note, “Cisco CRS-1 Carrier Routing System Gigabit Ethernet Physical Layer Interface Module Installation Note,” <http://www.cisco.com/en/US/docs/routers/crs/crs1/plim/installation/guide/eth6437.pdf>, February 2007.
- [31] Product Data Sheet, “CISCO CRS-1 Modular Services Card,” http://www.cisco.com/en/US/prod/collateral/routers/ps5763/ps5862/product_data_sheet09186a008022d5ee.pdf, November 2004.
- [32] A. Bianco, E. Bonetto, F. Musumeci, A. Pattavina, and M. Tornatore, “CapEx/OpEx evaluation of circuit vs packet switched optical networks,” in *Proc. of the ONDM, Brest, France*, April 2013.
- [33] E. Yetginer and G. N. Rouskas, “Power efficient traffic grooming in optical WDM networks,” in *Proc. of the GLOBECOM, Honolulu, USA*, December 2009.
- [34] G. Shen and R. S. Tucker, “Energy-minimized design for IP over WDM networks,” *Journal of Optical Communications and Networking*, vol. 1, no. 1, pp. 176–186, June 2009.
- [35] A. Coiro, M. Listanti, T. Squarcia, A. Valenti, and F. Matera, “Energy-minimised virtual topology design in IP over WDM backbone networks,” *IET Optoelectronics, Special Issue on Green Photonics*, vol. 6, no. 4, pp. 165–172, August 2012.
- [36] 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*, vol. 57, no. 9, pp. 2004–2020, June 2013.
- [37] A. Coiro, M. Polverini, A. Cianfrani, and M. Listanti, “Energy saving improvements in IP networks through table lookup bypass in router line cards,” in *Proc. of the ICNC, San Diego, USA*, January 2013.
- [38] A. Ahmad, A. Bianco, E. Bonetto, D. Cuda, G. Gavilanes Castillo, and F. Neri, “Power-aware logical topology design heuristics in wavelength-routing networks,” in *Proc. of the ONDM, Bologna, Italy*, February 2011.
- [39] B. Puype, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, “Multilayer traffic engineering for energy efficiency,” *Photonic Network Communications*, vol. 21, no. 2, pp. 127–140, April 2011.

- [40] ———, “Power reduction techniques in multilayer traffic engineering,” in *Proc. of the ICTON, Azores, Portugal*, June 2009.
- [41] B. Puype, D. Colle, M. Pickavet, and P. Demeester, “Energy efficient multilayer traffic engineering,” in *Proc. of the ECOC, Vienna, Austria*, September 2009.
- [42] F. Musumeci, M. Tornatore, J. López Vizcaíno, Y. Ye, and A. Pattavina, “Energy-efficiency of protected IP-over-WDM networks with sleep-mode devices,” *Journal of High Speed Networks*, vol. 1, no. 19, pp. 19–32, March 2013.
- [43] Y. Wu, L. Chiaraviglio, M. Mellia, and F. Neri, “Power-aware routing and wavelength assignment in optical networks,” in *Proc. of the ECOC, Vienna, Austria*, September 2009.
- [44] P. Wiatr, P. Monti, and L. Wosinska, “Power savings versus network performance in dynamically provisioned WDM networks,” *IEEE Communications Magazine*, vol. 50, no. 5, pp. 48–55, May 2012.
- [45] A. Jirattigalachote, Ç. Çavdar, P. Monti, L. Wosinska, and A. Tzanakaki, “Dynamic provisioning strategies for energy efficient WDM networks with dedicated path protection,” *Optical Switching and Networking, Special Issue on Green Communications*, vol. 8, no. 3, pp. 171–180, July 2011.
- [46] A. Muhammad, P. Monti, I. Cerutti, L. Wosinska, P. Castoldi, and A. Tzanakaki, “Energy-efficient WDM network planning with dedicated protection resources in sleep mode,” in *Proc. of the GLOBECOM, Miami, USA*, December 2010.
- [47] Ç. Çavdar, F. Buzluca, and L. Wosinska, “Energy-efficient design of survivable WDM networks with shared backup,” in *Proc. of the GLOBECOM, Miami, USA*, December 2010.
- [48] J. Perelló, A. Morea, S. Spadaro, and M. Tornatore, “Link vs. opto-electronic device sleep mode approaches in survivable green optical networks,” in *Proc. of the OFC, Anaheim, USA*, March 2013.
- [49] A. Morea, J. Perelló, and S. Spadaro, “Traffic variation-aware networking for energy efficient optical communications,” in *Proc. of the ONDM, Brest, France*, April 2013.
- [50] Ç. Çavdar, A. Yayimli, and L. Wosinska, “How to cut the electric bill by routing and wavelength assignment with time-zones and time-of-use prices,” in *Proc. of the ACP, Shanghai, China*, November 2011.
- [51] ———, “How to cut the electric bill in optical WDM networks with time-zones and time-of-use prices,” in *Proc. of the ECOC, Geneva, Switzerland*, September 2011.
- [52] A. Morea, G. Rizzelli, and M. Tornatore, “On the energy and cost trade-off of different energy-aware network design strategies,” in *Proc. of the OFC, Anaheim, USA*, March 2013.
- [53] G. Rizzelli, A. Morea, C. Dorize, O. Rival, and M. Tornatore, “On the energy impact of transmission reach for 100g IP-over-WDM translucent optical networks,” in *Proc. of the ECOC, Amsterdam, The Netherlands*, September 2012.

- [54] G. Rizzelli, A. Morea, and M. Tornatore, "An analysis of daily power consumption under different on-off IP-over-WDM translucent design approaches," in *Proc. of the OFC, Anaheim, USA*, March 2013.
- [55] G. Rizzelli, A. Morea, M. Tornatore, and A. Pattavina, "Reach-related energy consumption in IP-over-WDM 100g translucent networks," *Journal of Lightwave Technology*, vol. 31, no. 11, pp. 1828–1834, June 2013.
- [56] F. Farahmand, M. M. Hasan, I. Cerutti, J. P. Jue, and J. J. P. C. Rodrigues, "Differentiated energy savings in optical networks with grooming capabilities," in *Proc. of the GLOBECOM, Miami, USA*, December 2010.
- [57] M. M. Hasan, F. Farahmand, A. N. Patel, and J. P. Jue, "Traffic grooming in green optical networks," in *Proc. of the ICC, Cape Town, South Africa*, May 2010.
- [58] F. Farahmand, I. Cerutti, M. M. Hasan, and J. P. Jue, "Energy-efficiency of drop-and-continue traffic grooming," in *Proc. of the OFC, Los Angeles, USA*, March 2011.
- [59] A. Coiro, M. Listanti, and A. Valenti, "Impact of energy-aware topology design and adaptive routing at different layers in IP over WDM networks," in *Proc. of the Networks, Rome Italy*, October 2012.
- [60] L. Chiaraviglio, M. Mellia, and F. Neri, "Reducing power consumption in backbone networks," in *Proc. of the ICC, Dresden, Germany*, June 2009.
- [61] —, "Minimizing ISP network energy cost: Formulation and solutions," *Transactions on Networking*, vol. 20, no. 2, pp. 463–476, April 2012.
- [62] A. P. Bianzino, L. Chiaraviglio, and M. Mellia, "Distributed algorithms for green IP networks," in *Proc. of the INFOCOM workshop on Communications and Control for Sustainable Energy Systems, Orlando, USA*, March 2012.
- [63] H. Yonezu, G. Shan, S. Shimizu, D. Ishii, S. Okamoto, E. Oki, and N. Yamanaka, "Network power saving topology calculation method by power ing off links considering QoS," in *Proc. of the OECC, Sapporo, Japan*, July 2010.
- [64] H. Takeshita, N. Yamanaka, S. Okamoto, S. Shimizu, and G. Shan, "Energy efficient network design tool for green IP/Ethernet networks," *Optical Switching and Networking, Special Issue ONDM 2010*, vol. 9, no. 3, pp. 164–270, July 2012.
- [65] H. Yonezu, K. Kikuta, D. Ishii, S. Okamoto, E. Oki, and N. Yamanaka, "QoS aware energy optimal network topology design and dynamic link power management," in *Proc. of the ECOC, Torino, Italy*, September 2010.
- [66] B. Addis, A. Capone, G. Carello, L. Gianoli, and B. Sansò, "Energy aware management of resilient networks with shared protection," in *Proc. of the SustainIT, Pisa, Italy*, October 2012.

- [67] X. Dong, T. El-Gorashi, and J. M. H. Elmirghani, "IP over WDM networks employing renewable energy sources," *Journal of Lightwave Technology*, vol. 29, no. 1, pp. 3–14, January 2011.
- [68] X. Dong, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Energy efficient optical networks with minimized non-renewable power consumption," *Journal of Networks*, vol. 7, no. 5, pp. 821–831, May 2012.
- [69] X. Dong, T. El-Gorashi, and J. M. H. Elmirghani, "Renewable energy in IP over WDM networks," in *Proc. of the ICTON, Munich, Germany*, June 2010.
- [70] —, "Hybrid-power IP over WDM network," in *Proc. of WOCN, Colombo, Sri Lanka*, September 2010.
- [71] X. Dong, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "On the energy efficiency of physical topology design for IP over WDM networks," *Journal of Lightwave Technology*, vol. 30, no. 12, pp. 1931–1942, June 2012.
- [72] N. I. Osman, T. E. H. El-Gorashi, and J. M. H. Elmirghani, "Caching in green IP over WDM networks," *Journal of High Speed Networks*, vol. 1, no. 19, pp. 33–53, March 2013.
- [73] A. Bianco, E. Bonetto, D. Cuda, G. Gavilanes Castillo, M. Mellia, and F. Neri, "Energy-efficient design of wavelength-routing networks," in *Proc. of the ECOC, Torino, Italy*, September 2010.
- [74] P. Chowdhury, M. Tornatore, A. Nag, E. Ip, T. Wang, and B. Mukherjee, "On the design of energy-efficient Mixed-Line-Rate (MLR) optical networks," *Journal of Lightwave Technology*, vol. 30, no. 1, pp. 130–139, January 2012.
- [75] P. Chowdhury, M. Tornatore, and B. Mukherjee, "On the energy efficiency of mixed-line-rate networks," in *Proc. of the OFC, San Diego, USA*, March 2010.
- [76] B. Sansò and H. Mellah, "On reliability, performance and internet power consumption," in *Proc. of DRCN, Washington, UDA*, October 2009.
- [77] A. Silvestri, A. Valenti, S. Pompei, F. Matera, A. Cianfrani, and A. Coiro, "Energy saving in optical transport networks exploiting transmission properties and wavelength path optimization," *Optical Switching and Networking*, vol. 7, no. 3, pp. 108–114, July 2010.
- [78] A. Silvestri, A. Valenti, S. Pompei, F. Matera, and A. Cianfrani, "Wavelength path optimization in optical transport networks for energy saving," in *Proc. of the ICTON, Azores, Portugal*, June 2009.
- [79] M. Xia, M. Tornatore, Y. Zhang, P. Chowdhury, C. Martel, and B. Mukherjee, "Green provisioning for optical WDM networks," *Journal of Selected Topics in Quantum Electronics*, vol. 17, no. 2, pp. 437–445, March-April 2011.
- [80] —, "Greening the optical backbone network: A traffic engineering approach," in *Proc. of the ICC, Cape Town, South Africa*, May 2010.

-
- [81] R. M. Morais, J. Pedro, P. Monteiro, and A. N. Pinto, "Impact of node architecture in the power consumption and footprint requirements of optical transport networks," *Journal of Optical Communications and Networking*, vol. 5, no. 5, pp. 421–436, May 2013.
- [82] A. Klekamp, U. Gebhard, and F. Ilchmann, "Energy and cost efficiency of adaptive and Mixed-Line-Rate IP over DWDM networks," *Journal of Lightwave Technology*, vol. 30, no. 2, pp. 215–221, January 2012.
- [83] A. Nucci, A. Sridharan, and N. Taft, "The problem of synthetically generating IP traffic matrices: Initial recommendations," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 3, pp. 19–32, July 2005.
- [84] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, "An energy saving routing algorithm for a green OSPF protocol," in *Proc. of the INFOCOM Workshops, San Diego, USA*, March 2010.
- [85] A. Cianfrani, V. Eramo, M. Listanti, and M. Polverini, "An OSPF enhancement for energy saving in IP networks," in *Proc. of the INFOCOM Workshops, Shanghai, China*, April 2011.
- [86] F. Cuomo, A. Abbagnale, A. Cianfrani, and M. Polverini, "Keeping the connectivity and saving the energy in the Internet," in *Proc. of the INFOCOM Workshops, Shanghai, China*, April 2011.
- [87] F. Cuomo, A. Abbagnale, and M. Polverini, "ESOL: Energy saving in the Internet based on occurrence of links in routing paths," in *Proc. of the WoWMoM, Lucca, Italy*, June 2011.
- [88] F. Cuomo, A. Cianfrani, M. Polverini, and D. Mangione, "Network pruning for energy saving in the Internet," *Computer Networks*, vol. 56, no. 10, pp. 2355–2367, July 2012.
- [89] A. Cianfrani, V. Eramo, M. Listanti, M. Listanti, and A. V. Vasilakos, "An OSPF-integrated routing strategy for QoS-aware energy saving in IP backbone networks," *Transactions on Networking and Service Management*, vol. 9, no. 3, pp. 254–267, September 2012.
- [90] W. Shen, Y. Tsukishima, K. Yamada, and M. Jinno, "Power-efficient multi-layer traffic networking: Design and evaluation," in *Proc. of the ONDM, Kyoto, Japan*, February 2010.
- [91] A. Cianfrani, V. Eramo, M. Listanti, and M. Polverini, "Introducing routing standby in network nodes to improve energy savings techniques," in *Proc. of the e-Energy, Madrid, Spain*, May 2012.
- [92] A. Coiro, M. Listanti, A. Valenti, and F. Matera, "Reducing power consumption in wavelength routed networks by selective switch off of optical links," *Journal of Selected Topics in Quantum Electronics*, vol. 17, no. 2, pp. 428–436, March–April 2011.
- [93] A. Coiro, A. Valenti, F. Matera, and M. Settembre, "Reducing power consumption in core wavelength division multiplexing optical networks," *Fiber and Integrated Optics*, vol. 30, no. 3, pp. 166–177, June 2011.

- [94] A. Coiro, F. Iervini, and M. Listanti, “Distributed and adaptive interface switch off for internet energy saving,” in *Proc. of the International Conference on Computer Communications and Networks (ICCCN), New Delhi, India*, August 2011.
- [95] C. Lange and A. Gladisch, “Energy efficiency limits of load adaptive networks,” in *Proc. of the OFC, San Diego, USA*, March 2010.
- [96] —, “Limits of Energy Efficiency Improvements by Load-Adaptive Telecommunication Network Operation,” in *Proc. of the CTTE, Berlin, Germany*, May 2011.
- [97] Y.-W. Chen and C.-C. Chou, “Traffic modeling of a sub-network by using ARIMA,” in *Proc. of the ICII, Beijing, China*, October 2001.
- [98] X. Dong, A. Lawey, T. El-Gorashi, and J. M. H. Elmirghani, “Energy-Efficient Core Networks,” in *Proc. of the ONDM, Colchester, UK*, April 2012.
- [99] X. Dong, T. E. H. El-Gorashi, and J. M. H. Elmirghani, “Green IP over WDM networks with data centers,” *Journal of Lightwave Technology*, vol. 29, no. 12, pp. 1861–1880, June 2011.
- [100] S. Zhang, D. Shen, and C. Chan, “Energy efficient time-aware traffic grooming in wavelength routing networks,” in *Proc. of the GLOBECOM, Miami, USA*, December 2010.
- [101] B. G. Bathula, M. Alresheedi, and J. M. H. Elmirghani, “Energy efficient architectures for optical networks,” in *Proc. of the LCS, London, UK*, September 2009.
- [102] A. Coiro, M. Listanti, A. Valenti, and F. Matera, “Power-aware routing and wavelength assignment in multi-fiber optical networks,” *Journal of Optical Communications and Networking*, vol. 3, no. 11, pp. 816–829, November 2011.
- [103] A. Coiro, M. Listanti, and A. Valenti, “Dynamic power-aware routing and wavelength assignment for green WDM optical networks,” in *Proc. of the ICC, Kyoto, Japan*, June 2011.
- [104] W. Van Heddeghem, M. De Grotte, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, “Energy-efficiency in telecommunications networks: Link-by-link versus end-to-end grooming,” in *Proc. of the ONDM, Kyoto, Japan*, February 2010.
- [105] A. Tzanakaki, M. Anastasopoulos, K. Georgakilas, J. Buysse, M. De Leenheer, C. Develder, S. Peng, R. Nejabati, E. Escalona, D. Simeonidou, N. Ciulli, G. Landi, M. Brogle, A. Manfredi, E. López, J. Ferrer Riera, J. Garcia-Espin, P. Donaldio, G. Parladori, J. Jiménez, A. Tovar De Duenyas, P. Vicat-Blanc, J. Van der Ham, C. De Laat, M. Ghijsen, B. Belter, A. Binczewski, and M. Antoniak-Lewandowska, “Energy efficiency considerations in integrated IT and optical network resilient infrastructures,” in *Proc. of the ICTON, Stockholm, Sweden*, June 2011.
- [106] G. Rizzelli, A. Morea, M. Tornatore, and O. Rival, “Energy efficient traffic-aware design of on-off multi-layer translucent optical networks,” *Computer Networks*, vol. 56, no. 10, pp. 2443–2455, July 2012.

- [107] Y. Zhang, M. Tornatore, P. Chowdhury, and B. Mukherjee, “Energy optimization in IP-over-WDM networks,” *Optical Switching and Networking, Special Issue on Green Communications*, vol. 8, no. 3, pp. 171–180, July 2011.
- [108] M. Caria, M. Chamania, and A. Jukan, “To switch on or off: A simple case study on energy efficiency in IP-over-WDM networks,” in *Proc. of the HPSR, Cartagena, Spain*, July 2011.
- [109] ———, “A Comparative Performance Study of Load Adaptive Energy Saving Schemes for IP-Over-WDM Networks,” *Journal of Optical Communications and Networking*, vol. 4, no. 3, pp. 152–164, March 2012.
- [110] B. Mukherjee, D. Banerjee, S. Ramamurthy, and A. Mukherjee, “Some principles for designing a wide-area WDM optical network,” *Transactions on Networking*, vol. 4, no. 5, pp. 684–696, October 1996.
- [111] A. Betker, D. Kosiankowski, C. Lange, F. Pfeuffer, C. Raack, and A. Werner, “Energy efficiency in extensive multilayer core and regional networks with protection,” Konrad-Zuse-Zentrum für Informationstechnik Berlin, Tech. Rep. ZR 12-45, 2012.
- [112] 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,” *IEEE Communications Magazine*, vol. 51, no. 11, pp. 94–104, November 2013.
- [113] A. P. Bianzino, C. Chaudet, S. Moretti, J.-L. Rougier, L. Chiaraviglio, and E. Le Rouzic, “Enabling sleep mode in backbone IP-networks: a criticality-driven tradeoff,” in *Proc. of the ICC Workshop on Green Communications and Networking, Ottawa, Canada*, June 2012.
- [114] F. Idzikowski *et al.*, “Final report for the IRA “Energy-efficient use of network core resources,”” TREND Project, Deliverable D3.3, 2012.
- [115] “Rocketfuel: An ISP topology mapping engine,” <http://www.cs.washington.edu/research/networking/rocketfuel/>, April 2003.
- [116] M. Zhang, C. Yi, B. Liu, and B. Zhang, “GreenTE: Power-aware traffic engineering,” in *Proc. of the ICNP, Kyoto, Japan*, October 2010.
- [117] A. Tzanakaki, K. Katrinis, T. Politi, A. Stavdas, M. Pickavet, P. Van Daele, D. Simeonidou, M. J. O’Mahony, S. Aleksić, L. Wosinska, and P. Monti, “Power considerations towards a sustainable Pan-European network,” in *Proc. of the OFC, Los Angeles, USA*, March 2011.
- [118] ———, “Dimensioning the future Pan-European optical network with energy efficiency considerations,” *Journal of Optical Communications and Networking*, vol. 3, no. 4, pp. 272–280, April 2011.
- [119] A. Tzanakaki, M. Anastasopoulos, K. Georgakilas, J. Buysse, M. De Leenheer, C. Develder, S. Peng, R. Nejabati, E. Escalona, D. Simeonidou, N. Ciulli, G. Landi, M. Brogle, A. Manfredi, E. López, J. Riera, J. Garcia-Espin, P. Donadio, G. Parladori, and J. Jiménez, “Energy efficiency in integrated IT and optical network infrastructures: The GEYSERS approach,”

in *Proc. of the INFOCOM Workshop on Green Communications and Networking, Shanghai, China*, April 2011.

- [120] DICONET, “Definition of dynamic optical network architectures’,” DICONET Project, Deliverable D2.1, September 2010.
- [121] J. López Vizcaíno, Y. Ye, and I. Tafur Monroy, “Energy efficiency analysis for dynamic routing in optical transport networks,” in *Proc. of the ICC, Ottawa, Canada*, June 2012.
- [122] A. Coiro, A. Chiaraviglio, A. Cianfrani, M. Listanti, and M. Polverini, “Reducing power consumption in IP networks through table lookup bypass,” *Computer Networks*, vol. 64, no. 8 May 2014, pp. 125–142, May 2014.
- [123] L. Liu and B. Ramamurthy, “Rightsizing bundle link capacities for energy savings in the core network,” in *Proc. of the GLOBECOM, Houston, USA*, December 2011.
- [124] ———, “A dynamic local method for bandwidth adaptation in bundle links to conserve energy in core networks,” in *Proc. of the ANTS, Bangalore, India*, December 2011.
- [125] ———, “Investigation of energy savings in optical networks using bundle link rightsizing methods,” in *Proc. of the OFC, Anaheim, USA*, March 2013.
- [126] ———, “A dynamic local method for bandwidth adaptation in bundle links to conserve energy in core networks,” *Optical Switching and Networking*, vol. 10, no. 4, pp. 481–490, November 2013.
- [127] Y. Zhang, M. Tornatore, P. Chowdhury, and B. Mukherjee, “Time-aware energy conservation in IP-over-WDM networks,” in *Proc. of the PS, Monterey, USA*, July 2010.
- [128] A. Q. Lawey, T. E. H. El-Gorashi, and J. M. H. Elmirghani, “Distributed energy efficient clouds over core networks,” *Journal of Lightwave Technology*, vol. 32, no. 7, pp. 1261–1281, April 2014.
- [129] Y. Zhang, “6 months of abilene traffic matrices,” <http://www.cs.utexas.edu/~yzhang/research/AbileneTM/>, August 2009.
- [130] S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon, “Providing public intradomain traffic matrices to the research community,” *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 83–86, 2006.
- [131] “NOBEL – Next generation Optical networks for Broadband European Leadership,” <http://www.ist-nobel.org>, 2004–2007.
- [132] “Deutsche Forschungsnetz Verein,” <http://www.dfn.de>, DFN, 2010.
- [133] A. P. Bianzino, C. Chaudet, F. Larroca, R. D., and R. J.-L., “Energy-Aware Routing: a Reality Check,” in *Proc. of the GLOBECOM Workshops, Miami, USA*, December 2010.

- [134] F. Giroire, D. Mazaauric, J. Moulhierac, and B. Onfroy, “Minimizing routing energy consumption: from theoretical to practical results,” in *Proc. of GreenCom, Hangzhou, China*, December 2010.
- [135] ———, “Minimizing routing energy consumption: from theoretical to practical results,” Institut National de Recherche en Informatique et en Automatique (INRIA), Tech. Rep. 7234, July 2010.
- [136] 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, Special Issue on Green Communications*, vol. 8, no. 3, pp. 181–200, July 2011.
- [137] 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*, vol. 5, no. 5, pp. 394–410, May 2013.
- [138] A. Ahmad, A. Bianco, E. Bonetto, L. Chiaraviglio, and F. Idzikowski, “Energy-aware design of multilayer core networks,” *Journal of Optical Communications and Networking*, vol. 5, no. 10, pp. A127–A143, October 2013.
- [139] A. Bianco, E. Bonetto, and A. Ahmad, “Energy awareness in the design of optical core networks,” in *Proc. of the OFC, Anaheim, USA*, March 2013.
- [140] F. Idzikowski, L. Chiaraviglio, and F. Portoso, “Optimal design of green multi-layer core networks,” in *Proc. of the e-Energy, Madrid, Spain*, May 2012.
- [141] B. G. Bathula and J. M. H. Elmirghani, “Green Networks: Energy Efficient Design for Optical Networks,” in *Proc. of WOCN, Cairo, Egypt*, April 2009.
- [142] L. Wang, R. Lu, Q. Li, X. Zheng, and H. Zhang, “Energy efficient design for multi-shelf IP over WDM networks,” in *Proc. of the INFOCOM Workshop on Green Communications and Networking, Shanghai, China*, April 2011.
- [143] F. Musumeci, F. Vismara, V. Grkovic, M. Tornatore, and A. Pattavina, “On the energy efficiency of optical transport with time driven switching,” in *Proc. of the ICC, Kyoto, Japan*, June 2011.
- [144] E. Palkopoulou, D. A. Schupke, and T. Bauschert, “Energy efficiency and CAPEX minimization for backbone network planning: Is there a tradeoff?” in *Proc. of the ANTS, New Delhi, India*, December 2009.
- [145] DAUIN - Politecnico di Torino, “HPC Project,” <http://dauin-hpc.polito.it>, 2011.
- [146] R. Bolla, R. Bruschi, A. Cianfrani, and M. Listanti, “Introducing standby capabilities into next-generation network devices,” in *Proc. of PRESTO, Philadelphia, USA*, May 2010.
- [147] R. Bolla, R. Bruschi, A. Cianfrani, O. Jaramillo, and M. Listanti, “Energy-efficient sleeping modes for next-generation core networks,” in *Proc. of Future Network & Mobile Summit, Warsaw, Poland*, June 2011.

- [148] R. Bolla, R. Bruschi, A. Cianfrani, and M. Listanti, “Enabling backbone networks to sleep,” *IEEE Network*, vol. 25, no. 2, pp. 26–31, March/April 2011.
- [149] N. Yamanaka, S. Shimizu, and G. Shan, “Energy efficient network design tool for green IP/Ethernet networks,” in *Proc. of the ONDM, Kyoto, Japan*, February 2010.
- [150] A. P. Bianzino, L. Chiaraviglio, and M. Mellia, “GRiDA: a green distributed algorithm for backbone networks,” in *Proc. of the GreenCom, online*, September 2011.
- [151] A. P. Bianzino, L. Chiaraviglio, M. Mellia, and J.-L. Rougier, “GRiDA: GReen Distributed Algorithm for energy-efficient IP backbone networks,” *Computer Networks*, vol. 56, no. 14, pp. 3219–3232, September 2012.
- [152] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz, “Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios,” in *Proc. of the ONDM, Kyoto, Japan*, February 2010.
- [153] A. M. C. A. Koster, S. Orlowski, C. Raack, G. Baier, T. Engel, and P. Belotti, *Branch-and-cut techniques for solving realistic two-layer network design problems*. Springer Berlin Heidelberg, December 2009, ch. 3, pp. 95–118.
- [154] B. Fortz and M. Poss, “An improved benders decomposition applied to a multi-layer network design problem,” *Operations Research Letters*, vol. 37, pp. 359–364, September 2009.
- [155] G. Dahl, A. Martin, and M. Stoer, “Routing through virtual paths in layered telecommunication networks,” *Operations Research*, vol. 47, no. 5, pp. 693–702, October 1999.
- [156] S. Orlowski, “Optimal design of survivable multi-layer telecommunication networks,” PhD thesis, Technische Universität Berlin, May 2009.
- [157] E. Kubilinskas, “Design of multi-layer telecommunication networks,” PhD thesis, Lund University, February 2008.
- [158] P. Belotti, A. Capone, G. Carello, and F. Malucelli, “Multi-layer MPLS network design: The impact of statistical multiplexing,” *Computer Networks*, vol. 52, no. 6, pp. 1291–1307, April 2008.
- [159] P. Pavon-Marino, R. Aparicio-Pardo, B. Garcia-Manrubia, and N. Skorin-Kapov, “Virtual topology design and flow routing in optical networks under multi-hour traffic demand,” *Photonic Network Communications*, vol. 19, no. 1, pp. 42–54, February 2010.
- [160] A. Altin, H. Yaman, and M. C. Pinar, “The robust network loading problem under hose demand uncertainty: Formulation, polyhedral analysis, and computations,” *INFORMS Journal of Computing*, vol. Articles in Advance, 2010.
- [161] S. Mattia, “The robust network loading problem with dynamic routing,” Università di Roma la Sapienza, Tech. Rep. 3, 2010.
- [162] S. E. Terblanche, R. Wessälly, and J. M. Hattingh, “Solution strategies for the multi-hour network design problem,” in *Proc. of the INOC, Spa, Belgium*, April 2007.

- [163] S. Raghavan and D. Stanojević, “WDM optical design using branch-and-price,” April 2007, working paper, Robert H. Smith School of Business, University of Maryland.
- [164] A. Bley, U. Menne, R. Klähne, C. Raack, and R. Wessäly, “Multi-layer network design – a model-based optimization approach,” in *Proc. of the PGTS, Berlin, Germany*, 2008.
- [165] A. M. C. A. Koster, S. Orlowski, C. Raack, G. Baier, and T. Engel, *Single-layer Cuts for Multi-layer Network Design Problems*. Springer Verlag, February 2008, vol. 44, ch. 1, pp. 1–23.
- [166] R. Ahuja, T. Magnanti, and J. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [167] B. Fortz and M. Thorup, “Internet traffic engineering by optimizing OSPF weights,” in *Proc. of the INFOCOM, Tel-Aviv, Israel*, March 2000.
- [168] L. A. Wolsey and R. L. Rardin, “Valid inequalities and projecting the multicommodity extended formulation for uncapacitated fixed charge network flow problems,” *European Journal of Operational Research*, vol. 71, no. 1, pp. 95–109, November 1993.
- [169] D. Bienstock and O. Günlük, “Computational experience with a difficult mixed integer multicommodity flow problem,” *Mathematical Programming*, vol. 68, no. 1, pp. 213–237, 1995.
- [170] ———, “Capacitated network design – polyhedral structure and computation,” *Inform Journal on Computing*, vol. 8, pp. 243–259, 1996.
- [171] F. Ortega and L. A. Wolsey, “A branch-and-cut algorithm for the single-commodity, uncapacitated, fixed-charge network flow problem,” *Networks*, vol. 41, pp. 143–158, May 2003.
- [172] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W.H. Freeman and Company, San Francisco, 1979.
- [173] T. L. Magnanti, P. Mirchandani, and R. Vachani, “Modelling and solving the two-facility capacitated network loading problem,” *Operations Research*, vol. 43, pp. 142–157, 1995.
- [174] A. Dwivedi and R. E. Wagner, “Traffic model for USA long distance optimal network,” in *Proc. of the OFC, Baltimore, USA*, March 2000.
- [175] IBM – ILOG, “CPLEX,” <http://www.ilog.com/products/cplex/>, 2009.
- [176] F. Idzikowski, L. Chiaraviglio, and E. Bonetto, “EWA: an adaptive algorithm for energy saving in IP-over-WDM networks,” in *Proc. of the NOC, Vilanova i la Geltru, Spain*, June 2012.
- [177] L. Chiaraviglio, A. Cianfrani, A. Coiro, M. Listanti, J. Lorincz, and M. Polverini, “Increasing device lifetime in backbone networks with sleep modes,” in *Proc. of the SoftCOM, Primosten, Croatia*, September 2013.
- [178] A. Gençata and B. Mukherjee, “Virtual-topology adaptation for WDM mesh networks under dynamic traffic,” *Transactions on Networking*, vol. 11, no. 2, pp. 236–247, April 2003.

- [179] P. N. Tran and U. Killat, “Dynamic reconfiguration of logical topology for WDM networks under traffic changes,” in *Proc. of the IEEE Network Operations and Management Symposium (NOMS), Salvador, Brazil*, April 2008, pp. 279–286.
- [180] F. Idzikowski, E. Bonetto, and L. Chiaraviglio, “EWA – an adaptive algorithm using watermarks for energy saving in IP-over-WDM networks,” Technical University of Berlin, Telecommunication Networks Group, Tech. Rep. TKN-12-002, May 2012.
- [181] European Commission, “Eurostat database,” <http://epp.eurostat.ec.europa.eu/>, February 2012.
- [182] 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, Special Issue on Innovative Techniques for Power Consumption Saving in Telecommunication Networks*, vol. 2014, no. 806960, pp. 1–11, March 2014.
- [183] Juniper, “JUNOS Software Interfaces and Routing Configuration Guide,” <http://www.juniper.net/techpubs/software/junos-security/junos-security10.2/junos-security-swconfig-interfaces-and-routing/junos-security-swconfig-interfaces-and-routing.pdf>, June 2010.
- [184] “MiDORi Network Technologies Project (MiDORi),” <http://midori.yamanaka.ics.keio.ac.jp/>, July 2009.
- [185] N. Yamanaka, H. Takeshita, S. Okamoto, and S. Gao, “MiDORi: Energy efficient network based on optimizing network design tool, remote protocol and new layer-2 switch,” in *Proc. of the COIN, Jeju, South Korea*, July 2010.
- [186] H. Takeshita, Y. Oikawa, H. Yonezu, D. Ishii, S. Okamoto, and N. Yamanaka, “Demonstration of the self organized dynamic link power management by “MiDORi” energy optimal network topology design engine,” in *Proc. of the OFC, Los Angeles, USA*, March 2011.
- [187] S. Okamoto, Y. Nomura, H. Yonezu, H. Takeshita, and N. Yamanaka, “GMPLS-enabled, energy-efficient, self-organized network: MiDORi,” in *Proc. of the ACP, Shanghai, China*, November 2011.
- [188] Y. Nomura, H. Yonezu, D. Ishii, S. Okamoto, and N. Yamanaka, “Dynamic topology reconfiguration for energy efficient network with link power control: MiDORi,” in *Proc. of the WTC, Yokohama, Japan*, March 2012.
- [189] A. Hirao, Y. Nomura, H. Yonezu, H. Takeshita, D. Ishii, S. Okamoto, and N. Yamanaka, “Prim’s algorithm based P2MP energy-saving routing design for MiDORi,” in *Proc. of the COIN, Yokohama, Japan*, May 2012.
- [190] Y. Nomura, H. Yonezu, D. Ishii, S. Okamoto, and N. Yamanaka, “Dynamic topology reconfiguration for energy efficient multi-layer network using extended GMPLS with link power control,” in *Proc. of the OFC, Los Angeles, USA*, March 2012.

- [191] S. Okamoto, “Requirements of GMPLS extensions for energy efficient traffic engineering,” <http://tools.ietf.org/html/draft-okamoto-ccamp-midori-gmpls-extension-reqs-02>, Internet Engineering Task Force, March 2013.
- [192] D. Katz, D. M. Yeung, and K. Kompella, “Traffic engineering TE extensions for OSPF version 2,” <https://tools.ietf.org/html/rfc3630>, Internet Engineering Task Force, September 2003.
- [193] A. Valenti, A. Rufini, S. Pompei, F. Matera, S. Di Bartolo, C. Da Ponte, D. Del Buono, and G. Beleffi, “QoE and QoS comparison in an anycast digital television platform operating on passive optical network,” in *Proc. of the Networks, Rome, Italy*, October 2012.
- [194] L. Rea, S. Pompei, A. Valenti, F. Matera, C. Zema, and M. Settembre, “Quality of Service control based on Virtual Private Network services in a Wide Area Gigabit Ethernet optical test bed,” *Fiber and Integrated Optics*, vol. 27, no. 4, pp. 301–307, July 2008.
- [195] Cisco, “Per-Packet Load Balancing, IOS Release 12.2(28)SB Guide,” http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/pplb.pdf, February 2006.
- [196] Product Data Sheet, “Cisco XR 12000 and 12000 Series Gigabit Ethernet Line Cards,” http://www.cisco.com/en/US/prod/collateral/routers/ps6342/product_data_sheet0900aecd803f856f.pdf, January 2006.
- [197] Product Installation Note, “Cisco CRS-1 Carrier Routing System Packet-over-SONET/SDH Physical Layer Interface Module Installation Note,” <http://www.cisco.com/c/en/us/td/docs/routers/crs/crs1/plim/installation/guide/pos6436.pdf>, November 2006.
- [198] M. Mandviwalla and N. Tzeng, “Energy-efficient scheme for multiprocessor-based router linecards,” in *Proc. of the SAINT, Phoenix, USA*, January 2006.
- [199] N. McKeown, “Optics inside routers,” in *Proc. of the ECOC, Rimini, Italy*, September 2003.
- [200] Product Data Sheet, “T1600 Routing Node PIC Guide,” http://www.juniper.net/techpubs/en_US/release-independent/junos/information-products/topic-collections/hardware/t-series/t1600/pics/t1600-pic.pdf, not publicly available any more, April 2009.
- [201] —, “T640 Routing Node PIC Guide,” http://www.juniper.net/techpubs/en_US/release-independent/junos/information-products/topic-collections/hardware/t-series/t640/pics/t640-pic.pdf, not publicly available any more, April 2009.
- [202] —, “T320 Router PIC Guide,” http://www.juniper.net/techpubs/en_US/release-independent/junos/information-products/topic-collections/hardware/t-series/t320/pics/t320-pic.pdf, not publicly available any more, April 2009.
- [203] —, “Cisco 7500 Series Router,” https://www.cisco.com/en/US/prod/collateral/routers/ps359/product_data_sheet0900aecd800f5542.pdf, February 2005.
- [204] —, “Cisco 12000 Series Gigabit Switch Routers,” http://www.cisco.com/warp/public/cc/pd/rt/12000/prodlit/gsr_ov.pdf, not publicly available any more, December 1999.

- [205] J. Baliga, R. Ayre, K. Hinton, and R. S. Tucker, “Photonic switching and the energy bottleneck,” in *Proc. of the PS, San Francisco, USA*, August 2007.
- [206] J. Baliga, K. Hinton, and R. S. Tucker, “Energy consumption of the internet,” in *Proc. of the COIN-ACOFI, Melbourne, Australia*, June 2007.
- [207] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, “Energy consumption in optical IP networks,” *Journal of Lightwave Technology*, vol. 27, no. 13, pp. 2391–2403, July 2009.
- [208] Product Hardware Guide, “T320 Internet Routing Node,” <http://www.juniper.net/techpubs/hardware/t-series/t320/t320-hwguide/t320-hwguide.pdf>, June 2007.
- [209] —, “T640 Internet Routing Node,” <http://www.juniper.net/techpubs/hardware/t-series/t640/t640-hwguide/t640-hwguide.pdf>, February 2009.
- [210] —, “T1600 Internet Routing Node,” <http://www.juniper.net/techpubs/hardware/t-series/t1600/t1600-hwguide/t1600-hwguide.pdf>, April 2009.
- [211] —, “TX Matrix Platform,” <http://www.juniper.net/techpubs/hardware/t-series/tx-matrix/tx-matrix-hwguide/tx-matrix-hwguide.pdf>, December 2008.
- [212] Product Data Sheet, “2.5-GBPS Multirate Transponder Card for the CISCO ONS 15454 Multiservice Transport Platform,” http://www.cisco.com/application/pdf/en/us/guest/products/ps2006/c1650/ccmigration_09186a00801a24d9.pdf, not publicly available any more, September 2004.
- [213] —, “10-GBPS Multirate Transponder Card for the CISCO ONS 15454 Multiservice Transport Platform,” http://www.cisco.com/en/US/prod/collateral/optical/ps5724/ps2006/ps5320/product_data_sheet0900aecd80121bf7.pdf, June 2006.
- [214] —, “4 x 2.5-GBPS Muxponder Card for the CISCO ONS 15454 Multiservice Transport Platform,” http://www.cisco.com/en/US/prod/collateral/optical/ps5724/ps2006/ps5320/product_data_sheet0900aecd80121bee.pdf, June 2006.
- [215] —, “10-Gbps Full-Band Tunable Multirate Enhanced Transponder Card for the Cisco ONS 15454 Multiservice Transport Platform,” http://www.cisco.com/en/US/prod/collateral/optical/ps5724/ps2006/ps5320/product_data_sheet0900aecd803fc3e8.pdf, September 2008.
- [216] —, “Cisco 2-Channel SFP WDM Transponder,” https://www.cisco.com/en/US/prod/collateral/modules/ps5455/ps6575/product_data_sheet0900aecd80395a59.pdf, October 2007.
- [217] —, “Lucent Metropolis Multi-Wavelength Optical Amplifier,” http://www.alcatel-lucent.com/wps/DocumentStreamerServlet?LMSG_CABINET=Docs_and_Resource_Ctr&LMSG_CONTENT_FILE=Data_Sheets/Metropolis_Multi-Rate_Transponder_for_Service_Providers_Datasheet.pdf, 2004.
- [218] —, “Ciena F10-T 10G Transponder Module For the CN 4200 FlexSelect Advanced Services Platform Family,” http://www.ciena.com/files/F10-T_10G_Transponder_Module_A4_DS.pdf, not publicly available any more, 2008.

- [219] ———, “Fujitsu Flashwave 7200 - Transponder, Protection, Regenerator System,” <http://www.fujitsu.com/downloads/TEL/fnc/datasheets/flashwave7200.pdf>, 2002.
- [220] ———, “Fujitsu Flashwave 7300 - Transponder, Protection, Regenerator System,” <http://www.fujitsu.com/downloads/TEL/fnc/datasheets/flashwave7300.pdf>, 2002.
- [221] ———, “Fujitsu Flashwave 7500 small system (compact ROADM for Mid-Size Applications),” <http://www.fujitsu.com/downloads/TEL/fnc/datasheets/flashwave7500s.pdf>, October 2004.
- [222] ———, “Cisco Optical Regenerator OC-48/STM-16 Bidirectional Regenerator,” http://www.cisco.com/warp/public/cc/pd/olpl/metro/15104/prodlit/cor_ds.pdf, not publicly available any more, September 1999.
- [223] JDSU Data Sheet, “JDSU WaveShifter 200 ER 3R,” http://www.jdsu.com/productliterature/ws200er3r_ds.cms_tm_ae.pdf, August 2008.
- [224] Product Data Sheet, “iLynx STM-16 / OC-48 Regenerator,” http://www.cisco.com/en/US/prod/collateral/video/ps8806/ps8807/ps8847/product_data_sheet0900aecd806c4d3b.pdf, not publicly available any more, October 2003.
- [225] ———, “Fujitsu Flashwave 7700 Ultra Long Haul, High Capacity DWDM Solution,” <http://www.fujitsu.com/downloads/TEL/fnc/datasheets/flashwave7700.pdf>, 2002.
- [226] R. S. Tucker, R. Parthiban, J. Baliga, K. Hinton, R. W. A. Ayre, and W. V. Sorin, “Evolution of WDM optical IP networks: A cost and energy perspective,” *Journal of Lightwave Technology*, vol. 27, no. 3, pp. 243–252, February 2009.
- [227] Product Data Sheet, “CISCO ONS 15501 Erbium Doped Fiber Amplifier,” http://www.cisco.com/warp/public/cc/pd/olpl/metro/on15500/on15501/prodlit/ons15_ds.pdf, not publicly available any more, May 2003.
- [228] ———, “Fujitsu Flashwave 7600 Regional Long Haul DWDM Solution,” <http://www.fujitsu.com/downloads/TEL/fnc/datasheets/flashwave7600.pdf>, 2002.
- [229] ———, “Lucent Metropolis Multi-Wavelength Optical Amplifier,” [http://www1.alcatel-lucent.com/gsearch/accessFile.jhtml?sendURL=http%3A//www.alcatel-lucent.com/wps/DocumentStreamerServlet%3FLMSG_CABINET%3DDocs_and_Resource_Ctr%26LMSG_CONTENT_FILE%3DData_Sheets/Metropolis_Multi-Wavelength_Optical_Amplifier_Datasheet.pdf&sendCollection=RESOURCECENTER_DATA_SHEETS&sendTitle=Metropolis%AE+Multi-Wavelength+Optical+Amplifier+Datasheet+\(A4\)&specialColl=](http://www1.alcatel-lucent.com/gsearch/accessFile.jhtml?sendURL=http%3A//www.alcatel-lucent.com/wps/DocumentStreamerServlet%3FLMSG_CABINET%3DDocs_and_Resource_Ctr%26LMSG_CONTENT_FILE%3DData_Sheets/Metropolis_Multi-Wavelength_Optical_Amplifier_Datasheet.pdf&sendCollection=RESOURCECENTER_DATA_SHEETS&sendTitle=Metropolis%AE+Multi-Wavelength+Optical+Amplifier+Datasheet+(A4)&specialColl=), 2003.
- [230] ———, “JDSU OA 400 Amplifier Series,” http://www.jdsu.com/productliterature/oa400_ds.cms_ae_030706.pdf, 2006.
- [231] ———, “Infinera Optical Line Amplifier,” http://www.infinera.com/pdfs/ola/Infinera_OLA_Data_Sheet.pdf, January 2007.

- [232] —, “Ciena Configurable OADM for CoreStream Agility Optical Transport System,” http://www.ciena.com/files/Configurable_OADM_A4_DS.pdf, not publicly available any more, 2008.
- [233] —, “Fujitsu Flashwave 7500 Metro/Regional ROADM Platform,” <http://www.fujitsu.com/downloads/TEL/fnc/datasheets/flashwave7500.pdf>, September 2008.
- [234] Product brief, “Finisar Dynamic Wavelength Processor DWP 100 Wavelength Selective Switch,” http://www.finisar.com/sites/default/files/pdf/DWP%20100%20Product%20Brief%209_2011%20V6_0.pdf, September 2011.
- [235] —, “Finisar Dynamic Wavelength Processor DWP 50 Wavelength Selective Switch,” http://www.finisar.com/sites/default/files/pdf/DWP50_Wavelength_Selective_Switch_Product_Brief_9_2011_V3.pdf, September 2011.
- [236] O. Mäkinen, “Greener Networks to Exploit Multilayer Interworking,” in *Green Telco World Congress*, January 2009.
- [237] Product brochure, “Infinera DTN Switched WDM System,” http://www.infinera.com/pdfs/dtn/Infinera_DTN_Brochure.pdf, May 2007.
- [238] Product Data Sheet, “Calient DiamondWave FiberConnect - Fiber Optic Cross-Connection System,” http://www.calient.com/_docs/FiberConnect_Datasheet.pdf, not publicly available any more, 2008.
- [239] —, “Calient DiamondWave PXC - photonic switching system,” http://www.calient.com/_docs/pxc_datasheet.pdf, not publicly available any more, 2005.
- [240] —, “MRV OCC - Optical Cross Connect,” http://www.mrv.com/datasheets/MCC/PDF300/MRV-MCC-OCC_A4_HI.pdf, not publicly available any more, 2009.
- [241] —, “MRV OCC 96 - Optical Cross Connect,” http://www.mrv.com/sites/default/files/datasheets/us_pdfs/mrv-mcc-occ96_.pdf, September 2013.
- [242] Cisco Document ID: 47242, “Cisco 12000 Series Internet Router Architecture: Line Card Design,” <http://www.cisco.com/c/en/us/support/docs/routers/12000-series-routers/47242-arch12000-lcdesign.pdf>, July 2005.
- [243] E. Bonetto, L. Chiaraviglio, D. Cuda, F. Idzikowski, and F. Neri, “Exploiting traffic dynamics in power-aware logical topology design,” in *Proc. of the ECOC, Geneva, Switzerland*, September 2011.
- [244] F. Idzikowski, R. Duque, F. Jiménez, E. Le Rouzic, L. Chiaraviglio, and M. Ajmone Marsan, “Energy saving in optical operator networks: the challenges, the trend vision, and some results,” in *Proc. of the ECOC, Amsterdam, The Netherlands*, September 2012.