

POLITECNICO DI TORINO

SCUOLA DI DOTTORATO

Dottorato in Ingegneria Elettronica e delle Comunicazioni – XXV ciclo

Tesi di Dottorato

**Energy efficiency in wired  
telecommunication networks:  
approaches, methodologies and results**



**Edoardo BONETTO**

**Matr. 169643**

**Tutore**  
prof. Marco MELLIA

**Coordinatore del corso di dottorato**  
prof. Ivo MONTROSSET

Febbraio 2013



# Summary

This PhD thesis describes the research activities carried out with the Department of Electronics and Telecommunications of Politecnico di Torino. Most of the research activities have been developed inside the European project TREND (Towards Real Energy-efficient Network Design), a Network of Excellence financed by the European Union. In particular, collaborations have been conducted with Orange Labs, Networks and Carriers, of Lannion (France), Technische Universität Berlin (Germany) and FASTWEB (Italy).

The target of this thesis is the improvement of the energy efficiency in wired telecommunication networks. Indeed, several research studies have underlined that the already huge energy consumption is destined to increase in the next years since the traffic in telecommunication networks will continue to grow constantly due to the increasing popularity of new bandwidth-consuming applications, and the increase of population using them. This situation thus urges to find strategies that can mitigate the demand of energy in the networks.

The aim of the thesis is indeed to propose energy-aware strategies, that can improve the energy efficiency in the networks, and to evaluate which could be the real benefits of introducing these strategies. In particular, during the three years of research, proposals dedicated to the various network segments have been examined. Indeed, studies reported that each network segment is or will be responsible of the critical energy situation of the networks. For instance, the energy consumed in access and aggregation networks represents the largest contribute in the consumption of wired telecommunication networks and no reduction is expected in the future. While, in core networks, the growth of the consumption will be consistent, more than doubling its energy consumption quota in about 10 years.

Finally, attention has been devoted also to the performance and economical aspects. Indeed, the improvement of the energy efficiency may affect the performance perceived by the users and also it may require additional investments, which can be compensated by the economical savings obtained by reducing the energy bill.

# Contents

<b>Summary</b>	<b>iii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Energy consumption in wired telecommunication networks . . . . .	2
1.2 Energy efficiency: incentives and drawbacks . . . . .	3
1.3 In the thesis . . . . .	4
<b>2 Energy-aware Design in Core Networks</b>	<b>7</b>
2.1 The Logical Topology Design problem . . . . .	8
2.1.1 The Power-Aware Logical Topology Design . . . . .	8
2.1.2 Mixed Integer Linear Problem for the LTD . . . . .	9
2.2 Economical impact of the Power-Aware Logical Topology Design (LTD) (PA-LTD) . . . . .	11
2.3 Heuristics for the PA-LTD . . . . .	13
2.3.1 Greedy heuristic . . . . .	14
2.3.2 Meta-heuristic Techniques . . . . .	15
2.3.3 Performance Results . . . . .	16
2.4 Conclusion . . . . .	20
<b>3 Exploiting Traffic Dynamics for Energy-efficient Core Networks</b>	<b>21</b>
3.1 Multi-Period Power-Aware Logical Topology Design . . . . .	22
3.1.1 Network and node model . . . . .	22
3.1.2 Mathematical model . . . . .	23
3.1.3 Static Base Network (SBN) . . . . .	26
3.2 Heuristics . . . . .	26
3.2.1 Least Flow Algorithm (LFA) . . . . .	26
3.2.2 Genetic Algorithm (GA) . . . . .	27
3.3 Scenarios . . . . .	29
3.3.1 Networks and Traffic . . . . .	29
3.3.2 Power and Capital Expenditures (CapEx) Data . . . . .	29
3.4 Metrics . . . . .	30

3.5	Results . . . . .	31
3.5.1	Time Variation . . . . .	35
3.5.2	Impact of Static Base Network (SBN) . . . . .	37
3.5.3	Impact of Load Variation . . . . .	37
3.5.4	Power Breakdown . . . . .	39
3.5.5	Monetary Savings . . . . .	41
3.6	Discussion and Implementation Issues . . . . .	41
3.7	Conclusion . . . . .	43
<b>4</b>	<b>Switching Technologies in Core and Metro Networks</b>	<b>45</b>
4.1	Circuit Switching Technologies . . . . .	46
4.2	Time-Domain Optical Sub-Wavelength Switching Technologies . . . . .	46
4.2.1	Optical Burst Switching . . . . .	47
4.2.2	Time-domain Wavelength Interleaved Network . . . . .	48
4.2.3	Packet Optical Add Drop Multiplexer (POADM) . . . . .	49
4.3	Network Scenario . . . . .	49
4.4	Power Consumption Model . . . . .	50
4.5	Results and Discussion . . . . .	51
4.6	Conclusions . . . . .	53
<b>5</b>	<b>Energy Efficiency in Access and Aggregation Networks: Status and Perspectives</b>	<b>55</b>
5.1	Dataset . . . . .	56
5.1.1	Points of Presence (PoPs)' data . . . . .	56
5.1.2	ADSL lines' data . . . . .	57
5.2	Energy consumption characterization of the Points of Presence . . . . .	58
5.2.1	Overview of the main characteristics of the PoPs . . . . .	58
5.2.2	Energy consumption correlation with external temperature and traffic load . . . . .	59
5.2.3	Energy breakdown analysis . . . . .	64
5.2.4	Wrap-up . . . . .	65
5.3	Traffic characterization of Points of Presence and ADSL lines . . . . .	65
5.3.1	Macroscopic traffic analysis of the PoPs . . . . .	66
5.3.2	Microscopic traffic analysis of ADSL modems . . . . .	68
5.3.3	Wrap-up . . . . .	70
5.4	Energy efficient policies . . . . .	70
5.4.1	Energy proportionality in the PoPs . . . . .	70
5.4.2	Sleep mode policies at the ADSL modems . . . . .	73
5.5	Results . . . . .	76
5.5.1	Energy proportionality in the PoPs . . . . .	77
5.5.2	Sleep mode policies at the ADSL modems . . . . .	79

5.5.3 Achievable energy reduction and monetary savings . . . . .	84
5.6 Conclusions . . . . .	88
<b>6 Conclusions and Future Work</b>	<b>91</b>
<b>Bibliography</b>	<b>95</b>

# List of Tables

2.1	LT dependence with $\nu^O$ and $C^{TX}=500\$$ . . . . .	12
2.2	BEP in years . . . . .	13
3.1	Considered networks and traffic data . . . . .	30
3.2	Line card Daily Energy Consumption $E^{LC}$ [kWh] . . . . .	32
3.3	Total Daily Energy Consumption $E^{TOT}$ [kWh] . . . . .	33
3.4	Reconfiguration Ratio $\xi$ . . . . .	34
3.5	Overload Ratio $\phi$ . . . . .	35
3.6	Yearly Monetary Cost [k€] . . . . .	41
5.1	Summary of PoPs characteristics . . . . .	58
5.2	Energy breakdown of the PoPs . . . . .	64
5.3	Energy savings at the PoPs . . . . .	88

# List of Figures

- 1.1 Energy status in wired networks. Figure retrieved from the paper by C. Lange, D. Kosiankowski, C. Gerlach, F. Westphal, and A. Gladisch, “Energy Consumption of Telecommunication Networks”, published in the Proceedings of the ECOC 2009 [1]. . . . . 2
- 1.2 Traffic growth in backbone networks. Figure retrieved from the paper by E. Le Rouzic, “Network evolution and the impact in core networks”, published in the Proceedings of the ECOC 2010 [2]. . . . . 3
- 2.1 Average number of transceivers per node for uniform traffic cases in 16-nodes network . . . . . 17
- 2.2 Average number of transceivers per node in US backbone 24-nodes network traffic case . . . . . 17
- 2.3 Average number of hops per traffic request in US backbone 24-nodes network traffic case . . . . . 18
- 2.4  $\Delta$  power consumption percentage in the US backbone 24-nodes network traffic case . . . . . 18
- 3.1 Main idea of MP-PA-LTD. . . . . 22
- 3.2 Network model on exemplary logical and physical topologies. . . . . 23
- 3.3 Internet Protocol (IP) router model – configuration with 2 Line Card Shelves (LCSs) and a Fabric Card Shelf (FCS). . . . . 23
- 3.4 Power consumption of active Line Cards, total demand and total reconfigured traffic in the Abilene network. . . . . 36
- 3.5 Power consumption of active Line Cards, total demand and total reconfigured traffic in the Géant network. . . . . 36
- 3.6 SBN Variation: Energy consumption and reconfiguration ratio in the Abilene and Géant networks. . . . . 38
- 3.7 Load Variation: Energy consumption and reconfiguration ratio in the Abilene network. . . . . 39
- 3.8 Load Variation: Energy consumption and reconfiguration ratio in the Géant network. . . . . 39
- 3.9 Power Breakdown for the Abilene and Géant networks. . . . . 40
- 4.1 OBS functioning. . . . . 47



4.2	TWIN functioning. . . . .	48
4.3	POADM simplified structure and functioning. . . . .	49
4.4	Number of transmitters and receivers per node vs. flow per node . . . . .	51
4.5	Power consumption savings of sub-wavelength technologies with respect to translucent circuit switching . . . . .	52
5.1	Energy consumption and external temperature versus time. . . . .	60
5.2	Normalized energy consumption vs. external temperature for PoP <i>A</i> . . . . .	61
5.3	Energy-temperature correlation coefficient vs. free cooling capability. . . . .	61
5.4	Normalized energy consumption versus traffic load. . . . .	62
5.5	Energy consumption and normalized traffic profiles of PoP <i>B</i> during one day. . . . .	63
5.6	Energy consumption per user versus number of users. . . . .	64
5.7	Typical daily pattern at a PoP along one week. . . . .	66
5.8	Daily traffic distribution for a week of PoP <i>A</i> . . . . .	67
5.9	Traffic distribution of every PoP. . . . .	67
5.10	Examples of the evolution of the line rate for three different users. . . . .	68
5.11	CDF of the total activity of the monitored lines during the day. . . . .	69
5.12	CDF of the utilization of the monitored lines during the day. . . . .	69
5.13	Linear and DVFS energy-to-traffic proportional strategies . . . . .	71
5.14	Energy savings schemes . . . . .	74
5.15	Example of evolution of the line rate and the modem power consumption obtained using two sleep mode policies. . . . .	75
5.16	Energy savings for PoPs with linear proportional energy consumption . . . . .	77
5.17	Energy savings for PoPs with cubic proportional energy consumption . . . . .	77
5.18	Energy savings for PoPs with resource consolidation and two operating configurations . . . . .	78
5.19	Energy savings for PoPs with resource consolidation and three operating configurations . . . . .	78
5.20	Percentage of energy savings for PoP <i>B</i> with resource consolidation and three operating configurations . . . . .	79
5.21	CDF all $T_{off}$ periods for all the lines. Results refer to SSM with $MinRate = 1$ kb/s . . . . .	80
5.22	CDF of $T_{off-total}$ for all the lines. Results refer to SSM with $MinRate = 1$ kb/s. . . . .	81
5.23	Comparison of the average duration of total silent time $T_{off-total}$ using the SSM policy with different settings. . . . .	82
5.24	Comparison of the standard deviation of the duration of total silent time $T_{off-total}$ using the SSM policy with different settings. . . . .	82
5.25	Comparing the number of transitions $N_{trans}$ from sleep to ON state using the SSM policy with different settings. . . . .	83

5.26	CDF of the number of transitions $N_{trans}$ from sleep to ON state using the SSM policy with $MinRate = 1kb/s$ and different settings of $Idle$ . . . . .	84
5.27	Comparison of the average duration of total silent time $T_{off-total}$ using the 2SM policy with $Rate-low = 1 kb/s$ and $T_{active} = 2 min$ . . . . .	84
5.28	Comparing the number of state transitions $N_{trans}$ using the 2SM policy with $Rate-low = 1 kb/s$ and $T_{active} = 2 min$ . . . . .	85
5.29	Energy savings for the three models . . . . .	86
5.30	Average energy consumption per users. SSM with $MinRate = 1 kb/s$ and $T_{idle} = 15 min$ . 2SM with $Rate-low = 1 kb/s$ , $Rate-high = 100 kb/s$ , $Idle-short = 5 min$ , $Idle-long = 15 min$ and $T_{active} = 2 min$ . . . . .	86

# Chapter 1

## Introduction

Energy efficiency is one of the most challenging issues which science must face in the next future. Indeed, as the reduction of CO<sub>2</sub> emissions and the preservation of the environment against the climate change have started to be supported by specific policies of international organizations and governments, energy-efficiency has become a new driver for technological improvements across several research fields.

Information and Communication Technologies (ICT) in general is seen as a major player to reduce the energy consumption, due to the “move bit, not atoms” paradigm. However, ICT is also one of the major consumer of energy. Indeed, it has been estimated that ICT is responsible for a percentage ranging from 2% to 10% of the worldwide energy consumption, with the telecommunication networks that account for a large fraction of it [3,4].

In Italy, for example, Telecom Italia, the largest Italian Internet Service Provider (ISP), is the second largest consumer of electricity after the National Railway system, consuming more than 2 TWh per year [5], which is equivalent to the energy consumed by 660,000 families in one year. Similar and even more pessimistic considerations hold for the other developed countries (see for example [6] for the case of Japan).

For the future, trends predict a notable increase of the ICT power consumption up to 30% of the worldwide electricity demand by 2030 under a business-as-usual evaluation scenario [7]. Indeed, despite more energy efficient network devices are appearing on the market leading to a reduction of the joule/bit ratio, the increasing amount of bandwidth ensured to an always increasing amount of users makes the joule/user ratio dramatically increasing.

Not surprisingly, several initiatives already started that explicitly aim at “greening the Internet” [8–11]. Indeed, a more energy-conscious telecommunication network can significantly reduce global energy consumptions and costs.




Figure 1.1: Energy status in wired networks. Figure retrieved from the paper by C. Lange, D. Kosiankowski, C. Gerlach, F. Westphal, and A. Gladisch, “Energy Consumption of Telecommunication Networks”, published in the Proceedings of the ECOC 2009 [1].

## 1.1 Energy consumption in wired telecommunication networks

Before proposing and evaluating possible strategies that can improve the energy efficiency in wired telecommunication networks, we evaluate which is the energy consumption of the different network segments and which is their weight over the total consumed energy.

In Fig. 1.1, it is reported the relative energy consumed by each network segment for the year 2009 and the predicted trends for the following years. The total energy consumed by all the network segments has been normalized to 1 with respect to the year 2009. This analysis has been retrieved from [1].

It is possible to notice that the energy consumption in access networks accounts for the largest quota of the total consumed energy, about the 80%, in 2009. In the next years, this large amount is expected to slowly increase, since the energy consumed in access networks is proportional to the number of users which should not grow too much. Indeed, it has been evaluated that users’ terminals represent more than 65% of the overall energy consumed in access networks [12]. Moreover, there is a gradual shift from ADSL to PON access network technologies which have been evaluated as less energy consuming [13].

In the future, however, the consumption of access networks is going to decrease its weight over the total consumed energy. This is due to the fact that the consumption of the other network segments is dramatically increasing. In aggregation networks, indeed, there is a significant increase of the consumption, which in 2017 almost doubles with respect

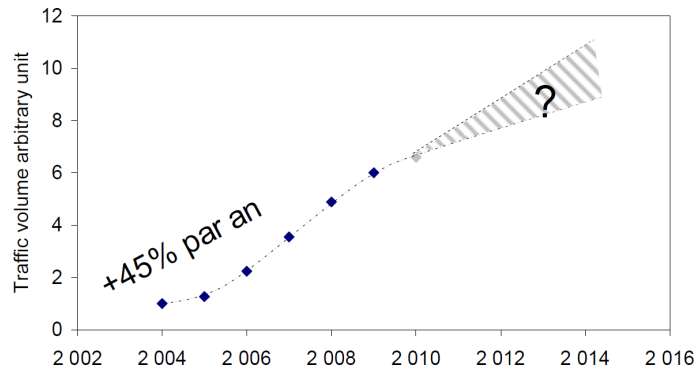


Figure 1.2: Traffic growth in backbone networks. Figure retrieved from the paper by E. Le Rouzic, “Network evolution and the impact in core networks”, published in the Proceedings of the ECOC 2010 [2].

to 2009. However, the main culprit of the energy increase are the core networks, whose consumption will explode and it will become the dominant part of the overall energy requirements.

The huge growth of the energy consumed by core networks is due to the increase of network traffic. Indeed, due to the fast spread of bandwidth-consuming applications, in the past years the traffic volume in core networks is increased with a growth rate of 45% per year [2] and it is expected that it will continue to increase even in the future years with a similar if not faster rate.

## 1.2 Energy efficiency: incentives and drawbacks

As already stated before, a motivation to improve the energy efficiency of telecommunication networks is the willingness to reduce the impact of the human activities over the environment. Nevertheless, other motivations are present that encourage to focus on the energy efficiency issue.

For instance, the next generation of networking devices has to increase the capability in order to sustain the increasing traffic demands, but this results in always larger and more power consuming devices, which imply also stricter cooling requirements. If this trend will continue, the scalability of devices, and thus of networks, may not be ensured since the cooling requirements may not be satisfied.

Other motivation to reduce the energy consumed in telecommunication network is represented by the possible economical savings. Indeed, as the energy costs are increasing, they are becoming a significant percentage of the Operative Expenditures (OpEx) of the ISPs, which willingly seek to reduce the increasing energy bill.

However, the improvement of the energy efficiency may bring also some drawbacks. Indeed, we can expect that in an energy efficient network the amount of resources is not

over-dimensioned. Thus, during the traffic peak hours, it may result into a degradation of the Quality of Service (QoS).

Furthermore, the implementation of energy efficient strategies in the devices and in the networks may require to develop new hardware, to modify the existing communications protocols or to change network configuration procedures. Each of these tasks requires additional work and it may increase the complexity with respect to the actual situation. Thus, before implementing these strategies it is better to evaluate their effectiveness and their impact over the QoS in order to understand if it is worth to implement them.

### 1.3 In the thesis

The analysis of the energy consumption in wired telecommunications networks shows that each segment is contributing or it will contribute in a significant way to the overall consumed energy. Thus, in this thesis, we propose energy saving strategies for each network segment in order to have a relevant impact over the global consumption of telecommunication networks.

In the following, we use real data to characterize the current situation and to evaluate the effectiveness of the proposed strategies. Combinatorial optimization together with heuristics have been the principal methods that we employ for the development of the strategies. Finally, the evaluation of the possible strategies is based on, apart from the energy saved, the QoS and the Capital Expenditures (CapEx), since we aim to propose strategies that can reduce the consumed energy, but in the meanwhile that do not impact negatively over other factors.

In the thesis, Chapter 2 focuses on the convenience of an energy-aware design of core networks. In details, it has been compared a design which minimizes the consumed energy with respect to one which minimizes the CapEx. The objective is to evaluate how much the minimization of the energy consumption affects the CapEx of the network. This problem has been tackled first formulating the problem using Mixed-Integer Linear Problem (MILP), and afterwards developing heuristics in order to scale also for large networks.

In Chapter 3, we evaluate the reduction of the energy consumption in core networks exploiting the dynamics of traffic in time. The aim of this activity is to evaluate the benefits of configuring dynamically a network according to the current traffic requirements. Indeed, networks are typically dimensioned and configured according to the traffic peaks and thus, in low traffic hours, devices are underutilized. Nevertheless, since no energy proportional technique is implemented in current devices, they consume always the same amount of energy. A multi-period configuration can then permit to reduce the energy consumption by switching off unused and/or underutilized resources. Only drawback is the need to reconfigure the devices between time periods which may impact on the Quality of Service provided. Thus, we develop some heuristic algorithms which solve the network

configuration trying to minimize the energy consumption and the reconfiguration cost.

Other possible solution for core networks is to use a more energy efficiency switching paradigm. In Chapter 4, we compare the energy efficiency of circuit switching and of time-domain optical sub-wavelength switching technologies. The main difference between these two switching paradigms is the granularity at which they can switch the traffic. While circuit switching technologies can switch just at a granularity equal to the wavelength, the other technologies can achieve a sub-wavelength switching granularity since they share the same wavelength in the time domain. Thus, in sub-wavelength switching technologies, the traffic grooming may be more effective allowing to reduce the number of network interfaces and of the energy consumption. Thus, we simulate the behavior of the different technologies in order to evaluate which technology promise to be the more energy efficient.

Energy consumption profiling and energy efficient policies for access and aggregation networks have been presented in Chapter 5. Focusing on access networks, we evaluate the possibility to implement sleep mode policies in ADSL modems at the DSLAM and users side. Since ADSL modems are consuming the same amount of energy independently from the traffic processed, strategies, that put the ADSL modems in a sleep mode state when no or little traffic is exchanged, have been proposed. A real traffic dataset provided by FASTWEB, one of the main network operator in Italy, users have been used to evaluate the ADSL lines activity and to evaluate the effectiveness of the proposed sleep mode policies. The collaboration with FASTWEB has also led to develop a research activity concerning the Points of Presence (PoPs), which are large switching nodes which act also as aggregation points for the end users. The objective of this activity is to profile the energy consumption of the PoPs and to determine which factors (such as the traffic processed or the external temperature at the PoP) impact over the energy consumption. This analysis have been performed considering data measured by FASTWEB at its PoPs. Afterwards, since it results that the energy consumed by the PoPs is not correlated to the traffic processed, strategies, aiming to adapt the energy consumption to the traffic load, have been examined and the achievable savings have been evaluated.

Lastly, conclusions and some discussions about future work are introduced in Chapter 6.





## Chapter 2

# Energy-aware Design in Core Networks

Nowadays, Wavelength Division Multiplexing (WDM) networks are the most common solution to design core networks because, exploiting the optical multiplicity in the transmission of several wavelength channels, they offer a large aggregate capacity.

In particular, WDM networks are configured according to the Wavelength Routed (WR) network paradigm. WR networks exploit WDM fiber links and Optical Cross-Connects (OXCs) to provide end-to-end optical circuits, called lightpaths, in which traffic is routed and it is transmitted transparently from the source to the destination node (i.e., without any electronic processing at intermediate nodes).

The WR design normally goes through two phases: the Logical Topology Design (LTD) and the Routing and Wavelength Assignment (RWA). In the LTD, given a node-to-node traffic matrix, the best set of lightpaths, which satisfies the traffic demands and optimizes a given target (e.g., the best set of lightpaths in terms of cost, price, and/or performance), is selected disregarding the physical topology. In the RWA, given the physical topology and the Logical Topology (LT), each lightpath is routed over the fiber links and a wavelength is assigned to it, so as to satisfy a given optimality target, subject to the wavelength uniqueness and, possibly, continuity constraint.

There are two main reasons why often the two problems are faced and solved independently, thereby taking a sub-optimal approach. First, the joint solution is often unfeasible because of computational complexity. Second, the owner of the physical infrastructure and the designer of the LT typically belong to independent organizations: a provider of transport capacity (often a telecommunication operator) faces the RWA problem on the basis of a LT defined by a provider of access services to users (an ISP or an Intranet administrator), which independently solves the LTD problem.

It was previously shown that Power-Aware RWA strategies can significantly reduce power consumption, even for small-size networks [14]. In the following, instead, we focus on the Power-Aware LTD (PA-LTD) problem: find the best energy-aware LT given a node-to-node traffic matrix.

Differently from [15], we do not assume any particular physical topology and node

architecture, but we investigate the LTD problem to find a good energy-wise balance between deployment of electronic and optical technologies. More precisely, we present a model to explore the cost trade-offs between transmitting data in the *optical* domain and switching/processing data in the *electronic* domain.

Thus, we propose a MILP formulation to find optimal power and cost-aware LTs. We analyze the economical impact of PA-LTD with respect to the case of designing the network using a Cost-Aware LTD (CA-LTD). We show that although energy-efficient telecommunication infrastructures can initially lead to larger CapEx, they finally translate into lower OpEx, thus making power-aware WR networks sustainable and more cost effective. In addition, due to the complexity of the MILP approach, we propose a greedy heuristic and a Genetic Algorithm (GA) which ensure performances close to the ones achieved by the MILP formulation.

The remaining of the chapter is organized as follows: in Sec. 2.1 we present the LTD problem and, in particular, we focus in Sec. 2.1.1 on its Power-Aware flavor and on the proposed MILP in Sec. 2.1.2. In Sec 2.2, we evaluate the effect on the CapEx and on the OpEx for the PA-LTD and the CA-LTD. The proposed heuristics are described in Sec. 2.3 and their performance is evaluated in Sec. 2.3.3. Conclusions are drawn in Sec. 2.4.

## 2.1 The Logical Topology Design problem

The LTD problem is agnostic of the network physical topology and it exploits only the knowledge of the traffic matrix, i.e. the amount of data that each pair of nodes in the network is willing to exchange. Indeed, the output of the LTD is the set of lightpaths which specifies a directed graph connecting nodes. Data can be sent from a source to a destination node either directly, in the optical domain, using one lightpath in a single-hop fashion, or following a multi-hop path in which more lightpaths are used. In the latter case, the traffic is switched and processed electronically between two lightpaths at each intermediate node of the multi-hop path. Electronic switching is assumed also due to buffering needs. Therefore, several traffic flows can share the same lightpath, possibly avoiding the deployment of new transmitter-receiver (TX-RX) pairs.

Classical optimization targets proposed for the LTD are the minimization of i) the link congestion, ii) the end-to-end latency or iii) the CapEx. In the following, we consider the LTD problem when power consumption minimization criteria are adopted.

### 2.1.1 The Power-Aware Logical Topology Design

We assume that the main sources of power consumption are i) optical transceivers (TX-RX pairs) and ii) electronic switching (performed at the intermediate and endpoint nodes). Transceivers perform the electronic-optical conversion (and vice versa) and electronic switching is needed to route through traffic and to perform grooming operations.

Therefore, minimizing the power consumption involves finding the best balance between usage of optical transmission and of electronic switching. Our intention is to trade off, from a power consumption perspective, the amount of electronic switching/processing with the number of optical TX-RX pairs which each node must be equipped with. Obviously, if the power consumed by optical TX-RX pairs is negligible with respect to the power required to switch data in electronics, an energy-efficient solution would lead to a fully connected topology, in which electronic switching is almost completely avoided. On the contrary, if the power consumed by the optical TX-RX pairs is considerably larger than the power needed to process data in the electronic domain, a topology with a small number of lightpaths like a star minimizes the power consumption.

However, the most energy-efficient topology is only part of the problem. It is also important to assess the economical impact of selecting one topology with respect to another. Coming back to the previous example, a full-mesh topology requires a number of transceivers scaling with  $O(N^2)$ , which implies a large CapEx investment, whereas the number of flows each node must process in electronics scales with  $O(N)$ , since each node processes only the self-generated traffic, typically leading to a lower OpEx. A star topology requires only  $O(N)$  transceivers, i.e., lower CapEx but, since the star hub needs to process a number of flows scaling with  $O(N^2)$  it might lead to higher OpEx.

### 2.1.2 Mixed Integer Linear Problem for the LTD

We model the PA-LTD and the CA-LTD via a MILP formulation. We consider a network with  $N$  nodes, where each node can be both source and destination of data traffic. The input parameter is the traffic matrix  $\underline{T} = [\lambda^{sd}]$ , where  $\lambda^{sd}$  specifies the traffic requirements from source node  $s$  to destination node  $d$  ( $s, d \in N$ ). Output variables are  $\lambda_{ij}^s$ , which indicate the traffic units of source  $s$  that travel on the lightpath starting from node  $i$  and ending in node  $j$  ( $s, i, j \in N$ ).

Since routing information is not an essential outcome for our purposes, we use an aggregated formulation which takes into account the overall traffic generated at each node and not each source-destination combination. Although the routing information is lost, the complexity of the MILP solution is reduced by one order of magnitude. Indeed, to solve the problem, we are interested in the number of TX-RX pairs required to transmit the traffic flowing between each pair of nodes and the aggregated amount of traffic each node is switching electronically.

Another set of output variables of the PA-LTD problem is  $n_{ij}$  which represents the number of lightpaths, i.e. the number of TX-RX pairs, needed between nodes  $i$  and  $j$  ( $i, j \in N$ ). The  $n_i$  represents the number of transmitters node  $i$  is equipped with and, also, the total number of receivers needed in all other nodes to receive traffic generated by  $i$ .

Another set of output variables of the PA-LTD problem is  $n_{ij}$  which represents the number of lightpaths, i.e. the number of TX-RX pairs, needed between nodes  $i$  and  $j$

( $i, j \in N$ ). The summation  $\sum_j n_{ij}$  represents the number of transmitters node  $i$  is equipped with and, also, the total number of receivers needed in all other nodes to receive traffic generated by  $i$ .

We assume synchronous transmission systems since they are the most commonly deployed today in core networks (consider for instance SONET or X-Gigabit Ethernet). Thus, the power consumed by TX-RX pairs, namely  $P^{TX}$ , does not depend on the transmission load, but only on their nominal bitrate  $B^{TX}$ .

Regarding the power consumption due to switching and grooming operations, we assumed that switching devices show a fixed and a variable power consumption contributions [16]. We assume that every node is equipped with the same type of switching device, thus, the constant contribution would be equal for all nodes and it does not add information to the LTD solution. The variable power consumption contribution is related to additional modules that the switching fabric has to be equipped with in order to sustain a given amount of traffic processed by a node. Thus, the power consumption depends on the amount of traffic switched at node  $i$  (namely  $\lambda_i$ ) [17]. More precisely, we assume that the power consumed by the switching fabric linearly depends on  $\lambda_i$  and can be computed as  $P^{SW}(\lambda_i) = \frac{\lambda_i}{\Delta}$ , where  $\Delta = \frac{B^{TX}}{P^{SW}(\lambda_i=B^{TX})}$  is a proportionality factor measured in  $bps/Watt$ , that indicates the amount of traffic switched by a switching fabric consuming 1 Watt.

To assess the energy trade off between optical transmission and electronic switching, we introduce  $\nu^O = \frac{P^{SW}(B^{TX})}{P^{TX}(B^{TX})}$  as the ratio between the power required to switch electronically a bandwidth equal to  $B^{TX}$  and the power consumption of a TX-RX pair (operating at bitrate  $B^{TX}$ ). Thus,  $\nu^O > 1$  ( $\nu^O < 1$ ) implies that transmitting data in the optical domain is power-wise convenient (not convenient) with respect to switching the same amount of information in the electronic domain. Finally, introducing  $\nu^O$ ,  $\Delta = \frac{B^{TX}}{\nu^O P^{TX}}$  and  $P^{SW} = \nu^O \frac{\lambda_i P^{TX}}{B^{TX}}$ . By varying  $\nu^O$ , we can trade the power efficiency relationship between optical and electronic technologies.

In the PA-LTD formulation the objective function is the minimization of the total power consumption

$$OF = P^O + P^E \quad (2.1)$$

in which, the first power contribute  $P^O$  represents the consumption of the optical TX-RX pairs and it is computed as  $P^O = P^{TX} \sum_{i,j} n_{ij}$ , and the second term  $P^E = \sum_i P^{SW}(\lambda_i)$  is the total power consumption due to switching data in electronics. The traffic each node must process is equal to  $\lambda_i = \sum_d \lambda^{id} + \sum_{j,s,i \neq s} \lambda_{ij}^s + \sum_s \lambda^{si}$ , equivalent respectively, to the sum of the total traffic produced by the node, the traffic being forwarded by the node and the traffic received by node  $i$ . Furthermore, the following constraints must be satisfied to return feasible LTs.

$$\sum_i (\lambda_{ij}^s - \lambda_{ji}^s) = \begin{cases} -\sum_d \lambda^{sd}, j = s, \forall (j,s) \\ \lambda^{sj}, j \neq s, \forall j \end{cases} \quad (2.2)$$

Eq. (2.2) expresses the flow conservation constraints which guarantees that traffic addressed to each node  $j$  is dropped while all the remaining traffic leaves node  $j$ .

$$\sum_s \lambda_{ij}^s \leq n_{ij} B^{TX}, \forall (i,j) \quad (2.3)$$

Eq. (2.3) ensures that the number of TX-RX pairs between nodes  $i$  and  $j$  satisfies the amount of traffic flowing between each pair of nodes.

$$\sum_{j,s} \lambda_{ij}^s \leq B^{SW}, \forall i \quad (2.4)$$

The maximum aggregate bandwidth that a node can electronically switch is limited to  $B^{SW}$  by Eq. (2.4).

$$\sum_j n_{ij} \leq \delta, \forall i ; \quad \sum_i n_{ij} \leq \delta, \forall j \quad (2.5)$$

Eqs. (2.5) fixe to  $\delta$  the maximum number of transceivers that each node can use respectively in transmission and reception.

For the CA-LTD, we assume that the monetary cost is dominated by monetary cost of the TX-RX pairs, since it is difficult to evaluate the monetary cost of switching data traffic in the electronic domain. Furthermore, we assume that each node is equipped with the same switching engine. Thus, this cost becomes a constant regardless of the chosen topology, and can be neglected. The CapEx is thus evaluated as  $C = C^{TX} \sum_{i,j} n_{ij}$ , where  $C^{TX}$  is the transceiver cost. In the CA-LTD thus we set  $OF = C$ .

## 2.2 Economical impact of the PA-LTD

We compare PA-LTD and CA-LTD both in terms of CapEx and OpEx, evaluating their Present Value of Annuity (PVA). In finance, the annuity refers to any terminating stream of fixed monetary flows over a specified period of time. In our case, the PVA is the sum of the capital costs  $C$  and of the operational costs  $O$ . Since we consider power as the main factor contributing to the OpEx [18], the yearly OpEx is  $O = P \cdot T \cdot c$ , where  $c$  is the monetary cost in \$(Wh) for a network operator and  $T$  is the number of hours in one year. Let  $k$  be the interest rate over the investment period. Thus,

$$\text{PVA}(n) = C + O \sum_{l=0}^n \frac{1}{(1+k)^l} \quad (2.6)$$

is the PVA at year  $n$ .

To study the impact of both PA-LT and CA-LT optimizations on OpEx and CapEx, we consider a small network of  $N = 16$  nodes, with an uniform traffic matrix, using the

Table 2.1: LT dependence with  $\nu^O$  and  $C^{TX}=500\$$ 

<i>low traffic</i>						
$\nu^O$	CapEx (k\$)		OpEx(k\$)		$\bar{n}$	
	PA	CA	PA	CA	PA	CA
1	15	15	1	1	1.8	1.8
2	15	15	1.6	1.6	1.8	1.8
5	15	15	3.3	3.3	1.8	1.8
10	15	15	6.2	6.2	1.8	1.8
20	120	15	11.4	12	15	1.8
30	120	15	15.5	17.8	15	1.8
<i>high traffic</i>						
$\nu^O$	CapEx (k\$)		OpEx(k\$)		$\bar{n}$	
	PA	CA	PA	CA	PA	CA
1	80	80	6.2	6.2	10	10
2	120	80	10.1	10.1	15	10
5	120	80	20.2	21.9	15	10
10	120	80	37	41.5	15	10
20	120	80	70.6	80.7	15	10
30	120	80	104.3	120	15	10

AMPL+CPLEX optimization environment. Among the several considered traffic scenarios, we show results for two network loads:  $\lambda^s = 0.9B^{TX}$  (*low traffic*) and  $\lambda^s = 7.5B^{TX}$  (*high traffic*). In our optimization runs,  $\nu^O$  ranges in [1, 20].

The transceiver characteristics are:  $B^{TX} = 10$  Gbit/s,  $P^{TX} = 8$  Watt. This power consumption behavior was derived from the analysis of separate transceiver blocks (laser, modulator, driver, TIA and clock and data recovery) from datasheets. For the cost of a transceiver at 10 Gbit/s, we assumed  $C^{TX} = [100,500]$  \$ [19] in order to represent the variability in the prices found, and the unitary cost  $c = 0.2$  \$/kWh [18]. In the evaluation of the PVA, the interest rate  $k$  is assumed to be 2%, according to the expected inflation rate.

The CapEx and the yearly OpEx for the PA-LT and the CA-LT are shown in Tab. 2.1 for the two traffic loads. The last two columns indicate the average number of lightpaths per node ( $\bar{n}$ ).

When the traffic is low, almost independently of  $\nu^O$ , both the PA-LTD and the CA-LTD return networks with a low connectivity degree (in particular, the outcome is a star with  $\bar{n} = (15 \times 1 + 1 \times 15)/16 = 1.8$ ). When optical transmission becomes highly convenient with respect to electronic switching ( $\nu^O \geq 20$ ), the PA-LT becomes a full mesh ( $\bar{n} = 15$ ), achieving a reduced yearly OpEx, which contributes to compensate for the higher CapEx. As the traffic load increases,  $\bar{n}$  increases too. Indeed, for the *high* traffic scenario and independently of  $\nu^O$ , CA-LTD optimizes the transceiver utilization, grooming traffic as much as possible, and the optimal solution becomes a partial mesh.

Table 2.2: BEP in years

$C^{TX}=500\$$						
$\nu^O$	1	2	5	10	20	30
<i>low</i>	-	-	-	-	-	104
<i>high</i>	-	-	32	10	4	3
$C^{TX}=100\$$						
<i>low</i>	-	-	-	-	60	10
<i>high</i>	-	-	5	2	1	0.5

On the contrary, since PA-LTD does not aim at maximizing transmitter utilization, PA-LTs usually show a higher  $\bar{n}$  which increases with  $\nu^O$ , converging to a full mesh already for  $\nu^O = 2$ .

Tab. 2.2 reports the Break Even Point (BEP), expressed in years, i.e., the time required to recover the larger CapEx for the PA-LT with respect to the CA-LT. A dash in the table means that PA-LTD and CA-LTD generate the same topology or that the resulting economical advantage is too small to provide a reasonable BEP.

As the load increases, it becomes convenient to move traffic from the electronic to the optical domain for smaller values of  $\nu^O$ , and the LT becomes increasingly connected.

For higher values of  $\nu^O$ , the BEP is shortened; it is also strongly influenced by the cost of the transceivers. These behaviors are confirmed on other traffic scenarios, not reported here for space limitations.

## 2.3 Heuristics for the PA-LTD

The PA-LTD, when properly formalized, may be solved, using any optimization software available on the market. These tools usually can find solutions with a certain degree of optimality, but the inherent complexity of the problem makes the time required to compute the solution impractical.

In the PA-LTD, the complexity scales with the size of the network, since the number of variables, required to mathematically describe the problem, is of the order of  $N^3$ . Indeed, the optimizer (AMPL+CPLEX software in our case) obtained solutions with 99% of optimality for a network with 24 nodes, but, as the number of nodes increases to be several tens, it became hard to get a good solution in a reasonable time. Note that, in real cases, networks comprise from several tens up to hundreds of nodes, making impossible to use any optimizer to solve the LTD problem.

Therefore, heuristic approaches are necessary since they may allow to find solutions sufficiently precise in a reasonable amount of time. For such reasons, we have developed an iterative greedy heuristic and a meta-heuristic based on a GA able to ensure good solutions in a limited amount of time.

### 2.3.1 Greedy heuristic

We firstly propose an iterative greedy heuristic, aiming to minimize the virtual topology power consumption by performing local optimal choices.

In the following descriptions, we set as target the power consumption of the network, but these algorithms can be easily adapted to several minimization problems, such as the capital expenditures minimization for instance.

All heuristics use a routing algorithm based on the Dijkstra algorithm and the routing information is at the base of the heuristic calculation. Since the routing is applied incrementally (that is, routing one traffic request at a time, using the current status of the network) the routing algorithm has to keep into account the current lightpath's traffic congestion to calculate further routes. Since the MILP formulation supports multi-path, i.e., traffic can be divided over severapaths can be used to go from a

Indeed, the algorithm routes a traffic request not on the shortest path, but on the shortest that supports the required capacity; so, before, starting the routing process each traffic demand is divided in  $h + 1$  different traffic requests such that  $\lambda^{sd} = (h + x) \times B^{TX}$ , where  $h = \lfloor \lambda^{sd} / B^{TX} \rfloor$  and  $x = \lambda^{sd} - B^{TX} \times h$ .

#### Less Energy Incremental Heuristic

---

##### Algorithm 1 LE-I heuristic

---

```

1: start from a void LT
2: for all traffic demands  $\lambda^{sd}$  do
3:   if not  $\exists$  a path with the required capacity between  $s$  and  $d$  then
4:     add a direct link between  $s$  and  $d$ 
5:   else
6:     compute the power consumption using the path with  $\lambda^{sd}$ 
7:     compute the power consumption of a new direct lightpath
8:     if power using path  $\leq$  power of using new lightpath then
9:       choose the path
10:    else
11:      add a direct lightpath between  $s$  and  $d$ 
12:    end if
13:  end if
14:  update lightpath traffic loads with this traffic demand
15: end for

```

---

The heuristic proposed is named Less Energy Incremental (LE-I) heuristic; Algorithm 1 describes how traffic lightpath between nodes are added into the networks and how flows are allocated on them. Starting from a void LT, each traffic demand is satisfied



by choosing the less power consuming alternative between adding a new direct link or by using a path already available, provided that there exist a path with enough bandwidth.

Note that the order in which the traffic demands are considered may affect the performance of the heuristic. Indeed, allocating first larger traffic requests, smaller traffic requests may use some spare capacity, while, when working in ascending order, it may be more difficult to allocate larger traffic requests on lightpaths which are partially used.

We consider three different orderings: increasing, decreasing and random order. In the first two cases, the traffic requests are ordered according to their traffic value in bps, from the smallest to the largest one and vice versa, respectively.

### 2.3.2 Meta-heuristic Techniques

Meta-heuristics are computational methods that implement a repeated search through a large set of possible solutions in order to optimize target function over the set. They make no assumptions over the problem to be solved, but make use of random mechanisms to avoid getting stuck in points or regions of the test set. Two examples of these techniques are the GAs (that uses natural evolution as principle) and simulating annealing (inspired in a metallurgical technique to reduce material defects); we propose a meta-heuristic based on a GA.

#### Brief Recall of Genetic Algorithms

The GA adopts elements typical of the natural evolution to converge to better solutions. It is based on a population of individuals, each one of them representing a solution of the problem, that evolves through generations and where only the more suitable individuals survive, i.e., pass to the next generations. Each individual represents a solution to the problem; an array is used to this represent the population. It may happen that some individuals do not represent any feasible solution. Therefore, before adding a new individual to the population, it is necessary to check for its feasibility. At the beginning of the GA, an initial population is created, generating new individuals from scratch. Each new individual is generated by setting every element of its array in a random value; if the new individual results feasible, it is added to the population; the process repeats until whole initial generation is obtained. With a suitable initial population, the evolution process begins. At each generation, a certain number of new individuals are obtained from the reproduction of individuals belonging to the current population. The reproduction consists in the following steps:

- *Selection*: selection of two individuals (the parents) from the previous generation (a certain priority criterion applies);
- *Crossover*: the parents' arrays are split into two sections of random size; then a new individual is obtained using a different section of each parent;

- *Mutation*: each element of the child solution array is mutated (changed) with a given probability;
- *Acceptance*: the feasibility of the individual is checked, if the resulting individual is unfeasible discard it, otherwise add the new individual to the population.

The reproduction phase repeats a given number of times, and at the end of the process, the population is constituted by old individuals and by their offspring. The population size has to be stable, only a part of the population can be selected for the next generation, so the GA defines a *fitness function* that identifies the more suitable individuals that should be selected to survive. The GA usually ends under certain convergence conditions: *i*) a maximum number of generations has been reached, *ii*) the best suited individual remains unchanged without improvement for more than a certain number of generations. The first case prevents the algorithm to run without control; the second one guarantees that a good individual is obtained, considering a priori that no further improvement is possible.

### Genetical Algorithm for PA-LTD

In our GA for PA-LTD, each element of the population represents a possible LT of our networks, thus, each element is represented through a set of lightpaths which represent a feasible solution (if the individual does not represent a suitable solution it is discarded immediately). We implement each individual through an array of size  $N^2$ , thus, elements  $l = \{0, 1, \dots, N^2 - 1\}$  of the array represent the number of transceivers used between node  $i$  and node  $j$ , with  $(i = \lfloor N^2/l \rfloor)$  and  $j = l \bmod N$ . Regarding the feasibility, a LT it is considered feasible only if it satisfies all the traffic requests, that is, all the traffic request can be routed over the current LT. We consider as fitness function the overall power consumption of the network; the lower is the consumption the more fit the individuals are. Therefore on each generation, only the less power consuming LTs survive to the next one. According to the classical methodology we consider constant population size of 30 LTs, reproducing an offspring of 20 LTs on each generation. Mutation occur such that each individual randomly change in average one element of the array, i.e., after reproduction, each element of each new generated individual is mutated with probability  $1/N^2$ . Finally, the algorithm was set to stop after  $N^2$  iterations without having a power improvement; this limit was set based on simulation experience.

### 2.3.3 Performance Results

The performance of the heuristics is evaluated comparing their results with optimal results obtained using an optimization software which in our case is AMPL+CPLEX. The heuristics are implemented using C programming language using an ad-hoc graph library for networking applications, which provides basic graph functionality and handling of network nodes and link status. After having analyzed several data sheets we assume to

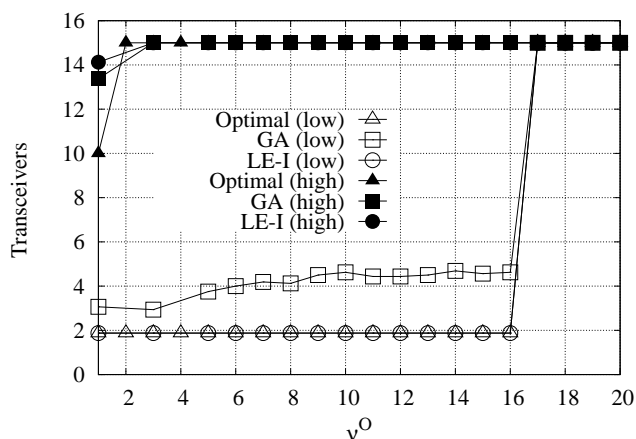


Figure 2.1: Average number of transceivers per node for uniform traffic cases in 16-nodes network




Figure 2.2: Average number of transceivers per node in US backbone 24-nodes network traffic case

use transceivers with capacity  $B^{TX} = 10$  Gbps and a nominal power consumption equal to  $P^{TX} = 8$  Watt [20]. We consider several optimization scenarios in which the parameter  $\nu^O$  assumes values in the range from 1 to 20, in order to test how the relative energy efficiency between optical and electronic technologies can affect the final result. The comparison is performed under an uniform traffic in a network of  $N = 16$  nodes under an unbalanced traffic matrix derived from the US backbone (retrieved from [21]). We use *rand* to indicate the random, *asc* for the ascending and *desc* for descending order in processing traffic requests.




Figure 2.3: Average number of hops per traffic request in US backbone 24-nodes network traffic case

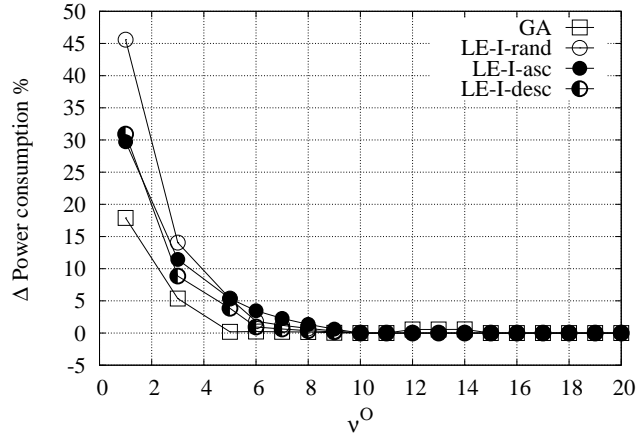


Figure 2.4:  $\Delta$  power consumption percentage in the US backbone 24-nodes network traffic case

For the uniform traffic case, we defined two scenarios: *low* traffic and *high* traffic. In the first scenario, all traffic demands are equal to  $\lambda^{sd} = 0.6 \text{ Gbps}$  ( $\forall s, d$ ), while in the second one they are set to  $\lambda^{sd} = 5 \text{ Gbps}$  ( $\forall s, d$ ). The first scenario represent the case in which all traffic flows of a node fits in a single lightpath, while the latter scenario is representative of the node to node traffic comparable to  $B^{TX}$ . Traffic requests are processed with no particular order since they are equal.

The Fig. 2.1 shows the average number of transceivers per node both in the *low* and *high* uniform traffic cases. Different algorithms are identified by different marker shapes.

In addition, white (black) filled markers identify the low (high) traffic.

Considering *low* uniform traffic, the optimal LT is a star for  $\nu^O < 16$  while it becomes a full mesh for larger  $\nu^O$ . LE-I heuristic results are optimal everywhere, while GA only approximates the star topology and converges to the optimal solution only when the optimal LT becomes a full mesh. The GA does not converge to this particular LT since it reaches the steady state before finding the optimal LT. In the high traffic the optimal solution is mesh topology (becoming a full mesh for  $\nu^O \geq 2$ ) and both algorithms catch this transition ensuring topology close to the optimal topology. Small differences are present for low value of  $\nu^O$  while coinciding topologies are found as  $\nu^O$  increases. Note that in the low uniform traffic scenario, there is just one transition from a low-connected to a highly-connected LT. Indeed, there exists a threshold equals to  $\bar{\nu}^0 = B^{TX}/\lambda^{sd}$  under the which is convenient to perform electronic switching, thus the star LT is the optimal solution, while above  $\bar{\nu}^0$  optical transmission becomes more energy efficient, hence leading to a full mesh LT.

The average number of transceivers per node of the 24 nodes US networks is depicted in Fig. 2.2. Again, as  $\nu^O$  increases the LT becomes more and more connected since optical transmission become more energy efficient with respect electronic switching. More in details, there is a transition toward a more connected LT for each value of  $\nu^O$  for which it becomes convenient for one or more traffic demands to be routed through a smaller number of hops. Regarding heuristics, also in this case performance are closer to the optimal one as  $\nu^O$  increases. Note that for low  $\nu^O$ , heuristics present some noise, i.e. they do not show a monotonic trend, and this is due to the fact that they get stuck to local minima. Note that almost independently from the traffic scenarios (more scenarios have been tested and results are not reported here due to space limits), the LE-I-*desc* tends to achieve a lower average number of transceivers than the other traffic orderings; processing the larger requests first through a single hop path allows to reuse some spare capacity to the smaller traffic requests which will be routed later.

Fig 2.3 shows the average number of hops as a function of  $\nu^O$ , as  $\nu^O$  increases and the LT becomes more connected, then, the average number of hops steadily decreases. CPLEX results are not available because of the aggregated formulation.

In addition, Fig. 2.4 shows the relative power consumption normalized to optimal LT, i.e.,  $\Delta P_{\%} = (OF^x - OF)/OF$ , where  $OF^x$  represents the total power consumption for the heuristics  $x$  ( $x \in \text{GA, LE-I-rand, LE-I-asc, LE-I-desc}$ ) evaluated as the summation of the power consumption due to electronic routing and optical transmission for the heuristic  $x$ . Finally,  $OF$  is evaluated according to Eq. (2.1). GA usually ensures lower power consumption than the LE-I heuristic, independently from the ordering used to accommodate traffic requests. Regarding the three different ordering methods of the LE-I, it seems that none of them presents a particular advantage, since the best performance ordering changes depending the value of  $\nu^O$ . However differences are negligible and it decreases as  $\nu^O$  increases.

## 2.4 Conclusion

We have explored the power awareness in the design of LTs in WR networks, the most common type of WDM core networks. We defined an optimal model (i.e., the MILP) and efficient heuristic algorithms. Moreover, we performed an economical analysis between the design of the LTs with power and cost minimization target.

We showed that power-aware approaches brings economical benefits and since we assumed worst-case prices (no economy of scale, no volume discounts), the economical advantages obtained by power-efficient techniques can be even more significant.

The MILP solved through commercial optimizers permit to infer the behavior of the LTs as the power cost between optics and electronics change. Significant differences in the LT design are highlighted that permit to exploit power benefits of the different technologies. Furthermore, the optimal solutions were taken as a benchmark for the performance of a heuristic and a meta-heuristic algorithms.

Using the heuristic techniques, it is possible to solve the LTD problem even for large scale networks, when a general purpose optimization software cannot be used because of the large execution times. We found that meta-heuristic approaches based on GA can lead to sub-optimal solutions, although close to optimal LTs.

## Chapter 3

# Exploiting Traffic Dynamics for Energy-efficient Core Networks

In this chapter, we focus on a multi-period approach for the design of LTs in WR networks. In particular, we exploit the traffic variability in order to configure the network in a more energy-efficient manner. Indeed, traffic is not constant over time, but changes significantly according to the hour (day/night). Indeed, instead of configuring the network according to the maximum requirements, it may be better to configure the network on a small time period scale. The intuition is that, configuring the network according to the current traffic requirements, it is possible to switch off unnecessary resources in low demand periods achieving energy savings.

However, a drawback of this approach is that the network has to be reconfigured between two subsequent time periods. Frequent and regular reconfigurations may cause a deterioration of the QoS perceived by the users, since the reconfiguration time is non-negligible. For this reason, we present a solution to reduce the power consumption with consideration of the amount of traffic that is reconfigured. In this way, while being energy-efficient, the negative effect of frequent reconfigurations is limited. In detail, we employ three heuristics to solve the problem, and compare them over an extensive set of scenarios, adopting as realistic assumptions on network topologies, traffic, CapEx and power values as possible.

The chapter is organized as follow. The general approach that we propose is presented in Section 3.1. The algorithms to solve the Multi-Period Power-Aware LTD (MP-PA-LTD) are described in Section 3.2. The considered scenarios are described in Section 3.3. Evaluation metrics are reported in Section 3.4. Section 3.5 reports the complete set of results. Discussion and implementation issues are reported in Section 3.6. Finally, conclusions are drawn in Section 3.7.




Figure 3.1: Main idea of MP-PA-LTD.

### 3.1 Multi-Period Power-Aware Logical Topology Design

In the MP-PA-LTD, we exploit the dynamic nature of traffic, which is not constant in time, but usually follows a day-night pattern with high traffic demands during the day and low traffic demands during the night. Today, network resources are allocated to constantly satisfy the maximum demands. These resources consume energy even if not used or underutilized. Indeed, power consumption in currently deployed network devices is basically independent of load [22]. Therefore, we divide the day into smaller time periods. Each period is characterized by a Traffic Matrix (TM) used to design the LT for the current period. Configuring the network for each of these periods allows utilization of resources to a higher extent, and consequently power can be saved by switching off some devices. Fig. 3.1 reports a schematic description of our solution. Normally, power consumption of the network is constant, independently of the current traffic (Fig. 3.1 left). With our approach, instead, the network power consumption is wisely adapted for each time period (Fig. 3.1 right). However, the network configuration (network nodes with installed devices in active or inactive states, established lightpaths, and routing of traffic demands) has to be changed between two consecutive time periods. Reconfiguring means adding or deleting lightpaths from the LT, which involves changing the routing of the traffic over the LT. In this research activity, we consider the latter as a metric to quantify the reconfiguration costs. Therefore, we define the reconfiguration cost for each lightpath as the amount of new traffic that has to be routed on the lightpath at the beginning of the new time period.

#### 3.1.1 Network and node model

We consider a backbone Internet Protocol (IP)-over-WDM network schematically modeled in Fig. 3.2. Fibers (blue solid lines) interconnect OXCs at the physical layer. They are used to realize the LT at the IP layer. We model the LT as a directed network composed of nodes and logical links (green dashed lines). A logical link consists of all the lightpaths established between a corresponding pair of IP routers, independently of their






Figure 3.2: Network model on exemplary logical and physical topologies.




Figure 3.3: IP router model – configuration with 2 LCSs and a FCS.

realization in the physical topology.

Each IP router has a modular structure, and is composed of one or more LCSs interconnected by FCSs, as depicted in Fig. 3.3. The LCS is the basic structure of the router. It is composed of the router chassis, switching fabric, cooling and power supply systems. It implements the control plane and data plane software. Furthermore, it is equipped with Line Cards (LCs) which are responsible for the physical connectivity of the router. Colored router interfaces are assumed. In the case that two or more LCSs are required, the FCSs are essential to interconnect the LCSs. Indeed, the FCSs contain the switching fabrics that are used for the interconnections of the LCSs.

In this work, we target reduction of the power by switching off LCs. This choice is due to the fact that LCs are expected to have significantly lower boot times than LCSs and FCSs, which makes their frequent activation and deactivation possible. We do not consider the power consumption of the OXCs and optical amplifiers since we focus on the LTD. Moreover, powering these devices on and off is expected to be time consuming.

### 3.1.2 Mathematical model

An informal description of the problem we target is the following. **Given** the current network configuration and the set of current traffic demands, **find** the set of lightpaths and corresponding LCs that must be powered on so that the weighted sum of the power

consumption and the reconfigured traffic is minimized, **subject to** flow conservation constraints, maximum lightpath utilization constraints, and constraints on the installed devices.

More formally, let us represent the LT of an IP-over-WDM network as a directed graph  $H = (V, L)$  where  $V$  is the set of all nodes in the network and  $L$  is the set of supplied logical links on which lightpaths can be established. Let  $L_{(i,j)} \subset L$  denote the set of supplied logical links between node  $i$  and node  $j$ . We exclude parallel logical links, therefore  $|L_{(i,j)}| = 1$  for each  $(i,j) \in V \times V$ .

Each logical link is formed by a trunk of parallel lightpaths. Let us define the capacity of a lightpath as  $C$ . Moreover, we denote by  $\delta \in (0,1]$  the maximum lightpath utilization. Each lightpath needs a transmitter and a receiver located in two LCs installed at the endpoints of the lightpath. The power consumption of a LC is denoted as  $\mathcal{P}^{LC}$ .

We consider a set of time periods  $T$  consisting of past and future time periods ( $T = T_{past} \cup T_{fut}$ ,  $T_{past} \cap T_{fut} = \emptyset$ ). We assume to know the traffic exchanged in the past time periods  $T_{past} \subset T$  (Fig. 3.1 left) and we determine the set of installed devices using this past traffic. We call this procedure Static Base Network (SBN) design. At the end of the SBN design, a set of installed devices is determined. Then we apply the MP-PA-LTD to the SBN on future time periods  $T_{fut} \subset T$  (Fig. 3.1 right). While time periods  $T_{past}$  regard the traffic before the SBN design, the future time periods  $T_{fut}$  concern future traffic which we assume to be unknown during the SBN design phase. The duration of each time period  $t \in T_{fut}$  is denoted as  $\Delta t$ . A Traffic Matrix (TM)  $D(t)$  for each time period  $t \in T$  contains traffic demands between the nodes  $(a,b) \in V \times V$  with values  $d^{ab}(t)$ . During the MP-PA-LTD constraints of the SBN are checked – the number of powered on devices at each node cannot exceed the number of installed devices.

We define the variables for each  $t \in T_{fut}$ . The flow variables  $f_{ij}^{ab}(t) \in \{0,1\}$  determine whether the traffic demand originated at node  $a$  and targeted to node  $b$  traverses the logical link from  $i$  to  $j$  at time  $t$ . Let us denote as  $r_{ij}^{ab}(t) \in \mathbb{R}_+$  the amount of reconfigured traffic (in Gbps) between  $a$  and  $b$  on the logical link from  $i$  to  $j$  at time  $t$ . The cost associated with each unit of reconfigured traffic is denoted by  $\mathcal{R}$ . The setting of  $\mathcal{R}$  allows trading between power minimization (low  $\mathcal{R}$ ) and low reconfigured traffic (high  $\mathcal{R}$ ). Thus, given a certain fixed value for  $\mathcal{R}$ , the solution of the problem represents the best trade-off between the consumed power and the reconfigured traffic. The value of  $\mathcal{R}$  has thus to be chosen according to the importance that the network operator gives either to power consumption or to reconfigured traffic. Moreover, let us introduce the variable  $y_l(t) \in \mathbb{Z}_+$  denoting the number of lightpaths established on the logical link  $l$  at time  $t$ . Finally, let  $x_i^{LC}(t) \in \mathbb{Z}_+$  be the number of LCs powered on at each node  $i$  at time  $t$ , which is bounded by the number of installed LCs  $X_i^{LC}$  in each node of the SBN.

Given the previous notations, we formalize the MILP for an arbitrary time period  $t$  as follows:

Objective:

$$\min \left( \mathcal{P}^{LC} \sum_{i \in V} x_i^{LC}(t) + \mathcal{R} \sum_{i \in V} \sum_{j \in V} \sum_{a \in V} \sum_{b \in V} r_{ij}^{ab}(t) \right) \quad (3.1)$$

Subject to:

$$\sum_{j \in V \setminus \{i\}} (f_{ij}^{ab}(t) - f_{ji}^{ab}(t)) = \begin{cases} 0 & i \neq a, i \neq b \\ 1 & i = a \\ -1 & i = b \end{cases}, \forall i, a, b \in V \quad (3.2)$$

$$d^{ab}(t) \cdot f_{ij}^{ab}(t) - d^{ab}(t-1) \cdot f_{ij}^{ab}(t-1) \leq r_{ij}^{ab}(t), \quad \forall i, j, a, b \in V \quad (3.3)$$

$$\sum_{a \in V} \sum_{b \in V} d^{ab}(t) \cdot f_{ij}^{ab}(t) \leq \delta \sum_{l \in L(i,j)} C \cdot y_l(t), \quad \forall i, j \in V \quad (3.4)$$

$$\sum_{j \in V} \sum_{l \in L(i,j)} y_l(t) \leq x_i^{LC}(t), \quad \forall i \in V \quad (3.5)$$

$$\sum_{i \in V} \sum_{l \in L(i,j)} y_l(t) \leq x_j^{LC}(t), \quad \forall j \in V \quad (3.6)$$

$$x_i^{LC}(t) \leq X_i^{LC}, \quad \forall i \in V \quad (3.7)$$

Control variables:  $f_{ij}^{ab}(t) \in \{0,1\}$ ,  $r_{ij}^{ab}(t) \in \mathbb{R}_+$ ,  $y_l(t) \in \mathbb{Z}_+$ ,  $x_i^{LC}(t) \in \mathbb{Z}_+$

The objective (3.1) minimizes the weighted sum of the total power consumed by LCs and of the total cost of reconfiguration. The cost of reconfiguration is the quantity of reconfigured traffic, weighted by the constant  $\mathcal{R}$ . Constraints (3.2) ensure the flow conservation of the traffic demands over the LT assuming unsplittable demands. Constraints (3.3) determine the reconfigured traffic between the current time period  $t$  and the preceding time period  $t-1$ .<sup>1</sup> Constraints (3.4) ensure enough bandwidth on a logical link to accommodate the traffic flows. Inequalities (3.5) and (3.6) compute the number of powered on LCs for each node. Finally, constraints (3.7) ensure that the number of LCs powered on is lower or equal than the number of installed LCs in the SBN.

<sup>1</sup>  $f_{ij}^{ab}(t-1)$  is computed in the previous time period and it is an input parameter for the current time period  $t$ .

### 3.1.3 Static Base Network (SBN)

The SBN is the network configuration used as a starting point for MP-PA-LTD. During this step, the sets of installed routers and of LCs are determined. In particular, we use the GA proposed in [23] to solve the problem with the objective of CapEx minimization.

The SBN is dimensioned to satisfy the maximum TM  $D_{SBN}$ , based on the set of past time periods  $T_{past}$ .

$$d_{SBN}^{ab} = \max_{t \in T_{past}} d^{ab}(t), \quad \forall a, b \in V \quad (3.8)$$

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 (LCs, LCSs and FCSs) may eventually sustain the future traffic.

The overprovisioning factor  $\gamma$  is defined as the ratio between the lightpath'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.

## 3.2 Heuristics

The presented MILP falls in the class of NP-hard problems, and therefore finding the optimal solution becomes impractical even for small networks. To overcome this issue, we follow a heuristic approach. In particular, we consider three different heuristics to solve the MP-PA-LTD: Least Flow Algorithm (LFA), GA and Energy Watermark Algorithm (EWA), with the last two being adaptive ones. Moreover, our algorithms assume shortest path routing and unsplittable traffic demands, two common assumptions usually adopted by network operators. Each algorithm returns for period  $t \in T_{fut}$  a network configuration defined as network nodes  $V$  with  $x_i^{LC}(t)$  powered on LCs (out of the installed  $X_i^{LC}$  LCs) established lightpaths forming logical links  $y_l(t)$  and IP routing of traffic demands  $f_{ij}^{ab}(t)$ . The rest of this section is devoted to the description of the algorithms LFA and GA, while a complete description of EWA is reported in [24, 25].

### 3.2.1 Least Flow Algorithm (LFA)

The LFA 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 proposed in [26], 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 of logical links from the SBN  $Y_l^{SBN}$ , the current TM  $D(t)$  and the maximum link utilization  $\delta$  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 TM cannot be satisfied by the SBN, LFA does not power off any link.

---

**Algorithm 2** Pseudo-code of LFA.

---

**Input:** Set  $Y_l^{SBN}$  from SBN, current TM  $D(t)$ , maximum utilization  $\delta$

**Output:** Updated netConfig

```

1: Lls=sortLeastFlow( $Y_l^{SBN}$ );
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

```

---

### 3.2.2 Genetic Algorithm (GA)

The GA, which has been introduced in Chapter 2, optimizes just the power consumption of the network. In the following, it has been adapted to solve the MP-PA-LTD to directly minimize the power consumption and the reconfiguration cost.

In particular, we define that an individual is feasible if its corresponding LT satisfies the TM and the constraint (3.7) on the maximum number of powered on devices. If the GA cannot find any feasible solution, constraint (3.7) is relaxed and dummy LCs are installed. However, the traffic routed over these additional LCs is considered as overload.

The fitness value is defined as the weighted sum of normalized power consumption and normalized reconfiguration cost. Thus, in this case, the GA optimizes the same objective (3.1) as the mathematical model presented in Section 3.1.2. The normalization of the power and of the reconfiguration cost is performed with respect to the possible worst case, i.e., every node has to process all the traffic of the current TM. The sum is weighted by the parameter  $\alpha \in [0,1]$ , which is used, similarly to the constant  $\mathcal{R}$  of (3.1), to trade between power and reconfiguration costs.

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 TM  $D(t)$ .

Moreover, three algorithm parameters have to be provided: the maximum number of generations without improvements  $M$ , the population size  $S$  and the offspring size  $K$ . 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 some other individuals that have been chosen as parents. 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).

---

**Algorithm 3** Pseudo-code of GA.

---

**Input:** netConfig from period  $t - 1$ , current TM  $D(t)$ ,  $\alpha$ ,  $M$ ,  $S$ ,  $K$

**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,  $\alpha$ ,  $K$ );
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);
```

---

### 3.3 Scenarios

We first describe the considered networks and the traffic assumptions. Next, we detail the CapEx and power consumption figures of the adopted network components.

#### 3.3.1 Networks and Traffic

We consider traffic demands that have been measured in three different networks: Abilene (12 nodes), Germany17 (17 nodes) and Géant (22 nodes). The physical supply networks and the corresponding TMs are available in [27, 28]. Moreover, each lightpath has capacity  $C$  equal to 40 Gbps.

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

Total period covered by  $T_{past}$  (design of the SBN) equals 1 month for Abilene and Géant. Note that fine granular TMs (5 min.) cover only one day in the Germany17 network [28]. Therefore we use 12 TMs with time granularity equal to 1 month over the year 2004 to dimension the SBN. Since these TMs 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 MP-PA-LTD algorithms under various traffic conditions.

The original TMs from [28] 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 3.1. Total demand (defined as  $\sum_{(a,b) \in V \times V} d^{ab}(t)$ ) over time is indicated in Figures 3.4(a) and 3.5(a) for two exemplary days (2004-08-27 for Abilene and 2005-06-11 for Géant). We do not present all traffic patterns due to lack of space. More details can be found in [24].

#### 3.3.2 Power and CapEx Data

Cisco CRS-1 router [29] was used to parameterize our study. Each LCS can accommodate up to  $W_{LCS} = 16$  LCs, and a FCS can interconnect up to  $W_{FCS} = 9$  LCSs. The power consumption values have been taken from [30]. In detail, a LCS consumes  $\mathcal{P}^{LCS} = 2920$  W

Table 3.1: Considered networks and traffic data [28]

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)

and a FCS consumes  $\mathcal{P}^{FCS} = 9100$  W. We consider as LC the “Cisco CRS 4-port 10-GE Tunable WDMPHY Interface Module” with the “Cisco CRS-1 Modular Services Card”, consuming together  $\mathcal{P}^{LC} = 500$  W.

The CapEx values used during the design of the SBN have been taken from [31]. In particular, we adopt values of 13.37, 16.67 and 53.35 units to LC, LCS and FCS, respectively.

### 3.4 Metrics

We describe the metrics adopted to assess the algorithms performance. We first introduce the power consumption of all active LCs in the network:

$$P^{LC}(t) = \mathcal{P}^{LC} \sum_{i \in V} x_i^{LC}(t) \quad (3.9)$$

The energy consumption due to LCs is then defined as:

$$E^{LC} = \sum_{t \in T_{fut}} P^{LC}(t) \Delta t \quad (3.10)$$

Moreover, we account also for the power of LCSs and FCSs. The number of LCSs used at each node is expressed as:

$$x_i^{LCS}(t) = \lceil x_i^{LC}(t) / W_{LCS} \rceil \quad (3.11)$$

Similarly, one FCS is used to interconnect up to  $W_{FCS}$  LCSs:

$$x_i^{FCS}(t) = \begin{cases} 0 & \text{if } x_i^{LCS}(t) \leq 1 \\ \lceil x_i^{LCS}(t) / W_{FCS} \rceil & \text{otherwise} \end{cases} \quad (3.12)$$

We define the power consumption of active LCSs in the network as:

$$P^{LCS}(t) = \mathcal{P}^{LCS} \sum_{i \in V} x_i^{LCS}(t) \quad (3.13)$$



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

$$P^{FCS}(t) = \mathcal{P}^{FCS} \sum_{i \in V} x_i^{FCS}(t) \quad (3.14)$$

The total network energy consumption is hence defined as:

$$E^{TOT} = \sum_{t \in T_{fut}} [P^{LC}(t) + P^{LCS}(t) + P^{FCS}(t)] \Delta t \quad (3.15)$$

Together with energy consumption, we are also interested in assessing the performance of our algorithms in terms of reconfigured traffic. We therefore introduce the reconfiguration ratio <sup>2</sup> over all subsequent pairs of time periods in  $T_{fut}$  as:

$$\xi = \frac{\sum_{t \in T_{fut}} \sum_{i \in V} \sum_{j \in V} \sum_{a \in V} \sum_{b \in V} r_{ij}^{ab}(t)}{\sum_{t \in T_{fut}} \sum_{a \in V} \sum_{b \in V} d^{ab}(t)} \quad (3.16)$$

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. Note that the reconfigured traffic may be counted multiple times (as many times as many logical links a traffic demand passes from its source to its target), so that  $\xi$  may be greater than 1.

Finally, we define the overload ratio metric  $\phi$  to capture overload traffic in all periods  $t \in T_{fut}$ :

$$\phi = \frac{\sum_{t \in T_{fut}} \sum_{i \in V} \sum_{j \in V} \chi_{ij}(t)}{\sum_{t \in T_{fut}} \sum_{a \in V} \sum_{b \in V} d^{ab}(t)} \quad (3.17)$$

where  $\chi_{ij}(t)$  is defined for each  $t \in T_{fut}$ ,  $i \in V$  and  $j \in V$  as:

$$\chi_{ij}(t) = \max \left( \sum_{a \in V, b \in V} d^{ab}(t) \cdot f_{ij}^{ab}(t) - \sum_{l \in L(i,j)} C y_l(t), 0 \right) \quad (3.18)$$

$\phi$  should be as low as possible (ideally 0) to prevent service disruptions and deterioration of QoS. Note that we account as overload each hop between a source and a destination, even though the traffic may get dropped already at the first hop. Therefore, also  $\phi$  may be greater than 1.

## 3.5 Results

We have implemented the proposed algorithms on custom simulators written in C++ and Java. The LFA and GA simulations are run on a server machine equipped with 2 Quad

<sup>2</sup>Note that  $r_{ij}^{ab}(t)$  and  $\xi$  take no value in the first period of  $T_{fut}$ .

Table 3.2: 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

CPU at 2.66 GHz and 4 GB of RAM. The EWA simulations are run on a personal computer with a Dual Core CPU at 2.4 GHz and 2 GB of RAM. In particular, we start considering SBN with  $d_{|V|} = 300$  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  $K = 20$  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: LCs daily energy consumption  $E^{LC}$ , total daily energy consumption  $E^{TOT}$ , reconfiguration ratio  $\xi$ , and overload ratio  $\phi$ . Table 3.2 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 LCs 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  $\gamma$ . This is caused by the fact that the traffic data available for the Germany17 network contains only one day of TMs with fine granularity (see Section 3.3.1 and [28]). Therefore the TMs of granularity of one month were used to design the SBN (see Table 3.1). Since the values in the TMs 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 TMs (see Eq. 3.8). 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 [27]. 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 the whole logical links, however since most of them 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 3.2).

Table 3.3: 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

In the following, we compare the heuristics considering the total energy consumption  $E^{TOT}$ . Table 3.3 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 LCs 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.

As the next step, we compare the heuristics in terms of the reconfiguration ratio  $\xi$ . Table 3.4 reports the results for the different scenarios. Intuitively, a ratio  $\xi$  approaching 1 indicates that the amount of reconfigured traffic is comparable with the total amount of traffic demands. This 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,  $\xi$  is typically larger than 0.66 for LFA. On the contrary,  $\xi$  is clearly reduced by the other heuristics, being lower than 0.15 and 0.18 for EWA and GA, respectively. This means that targeting explicitly the reconfiguration cost has a positive feedback on  $\xi$ . Similar considerations hold for the Germany17 and Géant network. Finally, observe also that  $\xi$  does not significantly change over the different days, i.e., between working days and weekends. This is due to the fact that a larger amount of traffic exchanged in the network (denominator of Eq. 3.16) implies also a larger amount of reconfigured traffic (numerator of Eq. 3.16).

Table 3.4: Reconfiguration Ratio  $\xi$

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

Finally, Table 3.5 reports the overload ratio over the different scenarios. Intuitively,  $\phi$  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.<sup>3</sup> 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

<sup>3</sup>Recall that we have adopted different time periods for designing the SBN and for evaluating the heuristics.

they integrate the possibility to add lightpaths while satisfying constraint of installed LCs in each network node. Table 3.5 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 LCs. This is due to the fact that the SBN is designed with  $d_{SBN}^{ab}$  calculated using monthly averages of demand values (see Section 3.3). Finally, the Table 3.5 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. Finally, GA avoids overload for all the periods in the Géant network.

Table 3.5: Overload Ratio  $\phi$ 

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

### 3.5.1 Time Variation

We then investigate the temporal behavior of the algorithms. In particular, we start considering the Abilene network<sup>4</sup> and the day 2004-08-27. Fig. 3.4(a) shows the power consumption of LCs versus time  $P^{LC}(t)$ . Power follows a day-night trend of traffic (right y-axis) 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 TM. This in turn leads to a new, less aggressive configuration of logical

<sup>4</sup>We do not change the time zone of the original TMs [28], therefore the day-night patterns in Abilene look shifted in time.

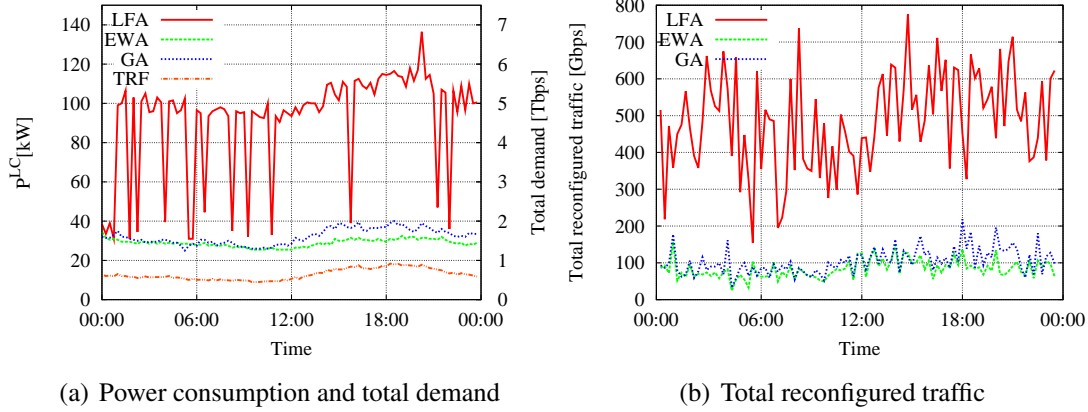


Figure 3.4: Power consumption of active Line Cards, total demand and total reconfigured traffic in the Abilene network.

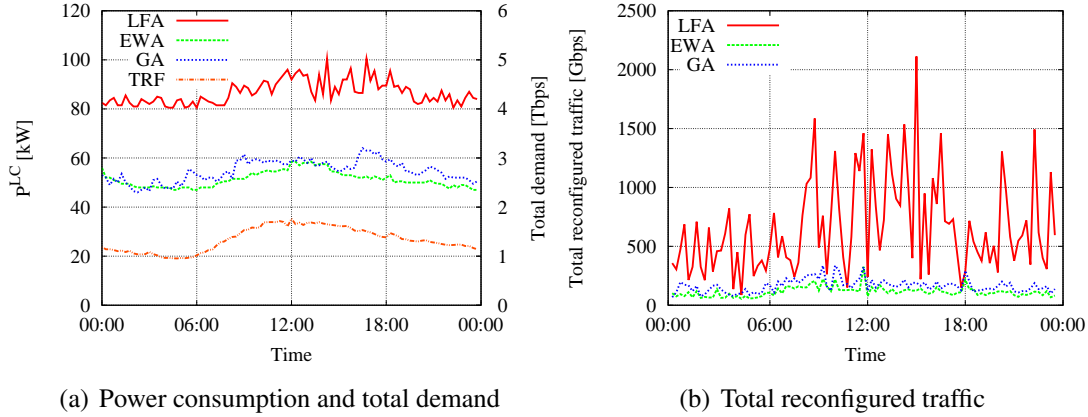


Figure 3.5: Power consumption of active Line Cards, total demand and total reconfigured traffic in the Géant network.

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 LCs installed in the SBN is more than 142 kW (not reported in the figure), thus we can conclude that our heuristics are effective in reducing the power consumption. To give more insight, Fig. 3.4(b) reports the total reconfigured traffic (sum of  $r_{ij}^{ab}(t)$  in (3.1)), 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 2005-06-10 day.

Fig. 3.5(a) reports power consumption of LCs versus time  $P^{LC}(t)$ . In this case, the power consumption of the always on network (SBN) is 276 kW, while our heuristics guarantee power consumption always lower than 105 kW. Again, a day-night behavior clearly emerges. Finally, the total reconfigured traffic is reported in 3.5(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.

### 3.5.2 Impact of Static Base Network (SBN)

In this section, we investigate the impact of the SBN design on the performance of our heuristics. In particular, we vary the overprovisioning factor  $\gamma$  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 day 2004-08-27 for Abilene, 2005-06-10 for Géant and 2005-02-15 for Germany17.

Fig. 3.6 reports  $E^{LC}$  and  $\xi$  for the different scenarios.  $\gamma$  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 LCs 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  $\xi$ , 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 LCs 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  $\gamma$ . This is due to the fact that both algorithms are not restricted to logical links existing in the SBN, and that they consider establishment and release of single lightpaths, and not whole logical links.

### 3.5.3 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  $\gamma$  set to 0.5. We consider the day 2004-08-27 of Abilene, and the day 2005-06-10 for Géant. Fig. 3.7 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 LCs has to be always powered on to guarantee connectivity. However, as traffic increases LFA consumes more energy than GA

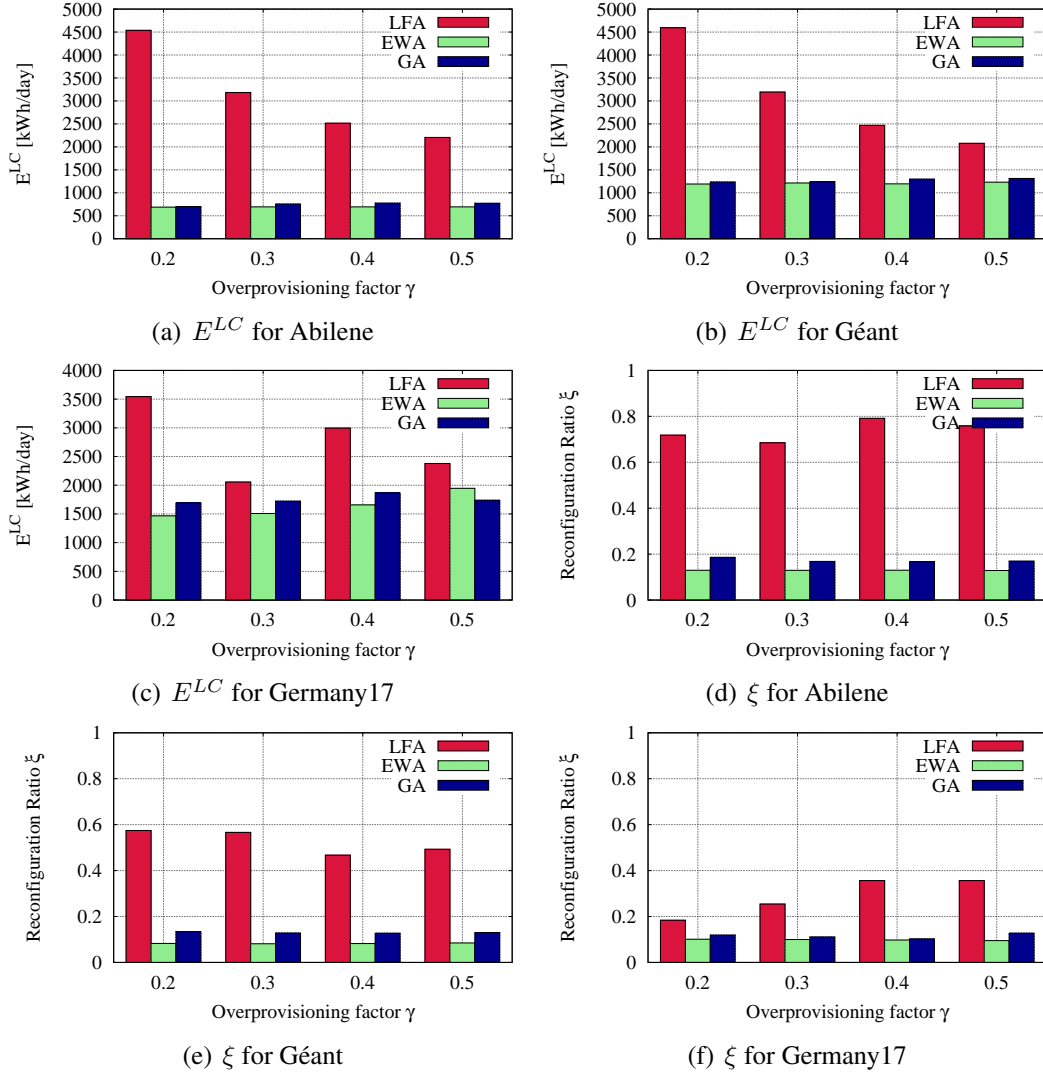


Figure 3.6: SBN Variation: Energy consumption and reconfiguration ratio in the Abilene and Géant networks.

and EWA. Moreover,  $\xi$  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. 3.8. It is interesting to note that the GA consumes slightly more energy than LFA under low load ( $d_{|V|} = 100$  Gpn).






Figure 3.7: Load Variation: Energy consumption and reconfiguration ratio in the Abilene network.




Figure 3.8: Load Variation: Energy consumption and reconfiguration ratio in the Géant network.

### 3.5.4 Power Breakdown

We then investigate how much energy is consumed by components of each type, differentiating between LCs, LCSs, and FCSs. In particular, the minimization of power due to LCs is targeted by the heuristics, while the power due to LCSs and FCSs is computed in a post-processing phase. Fig. 3.9 on the top reports the power consumption versus time for the Abilene over 2004-08-27. The LFA heuristic (Fig. 3.9(a)) requires all types of components to be powered on over the whole day, with a large amount of power consumed by LCs during the day. The EWA heuristic (Fig. 3.9(b)) instead does not require any FCS, and a lower amount of power is consumed by LCs 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. 3.9(c) is similar to EWA, with higher variation of power consumed by LCs.

We then consider the Géant network over 2005-06-10 (the second row in Fig. 3.9). Interestingly, all the algorithms require to use FCSs. This is due to the fact that in this

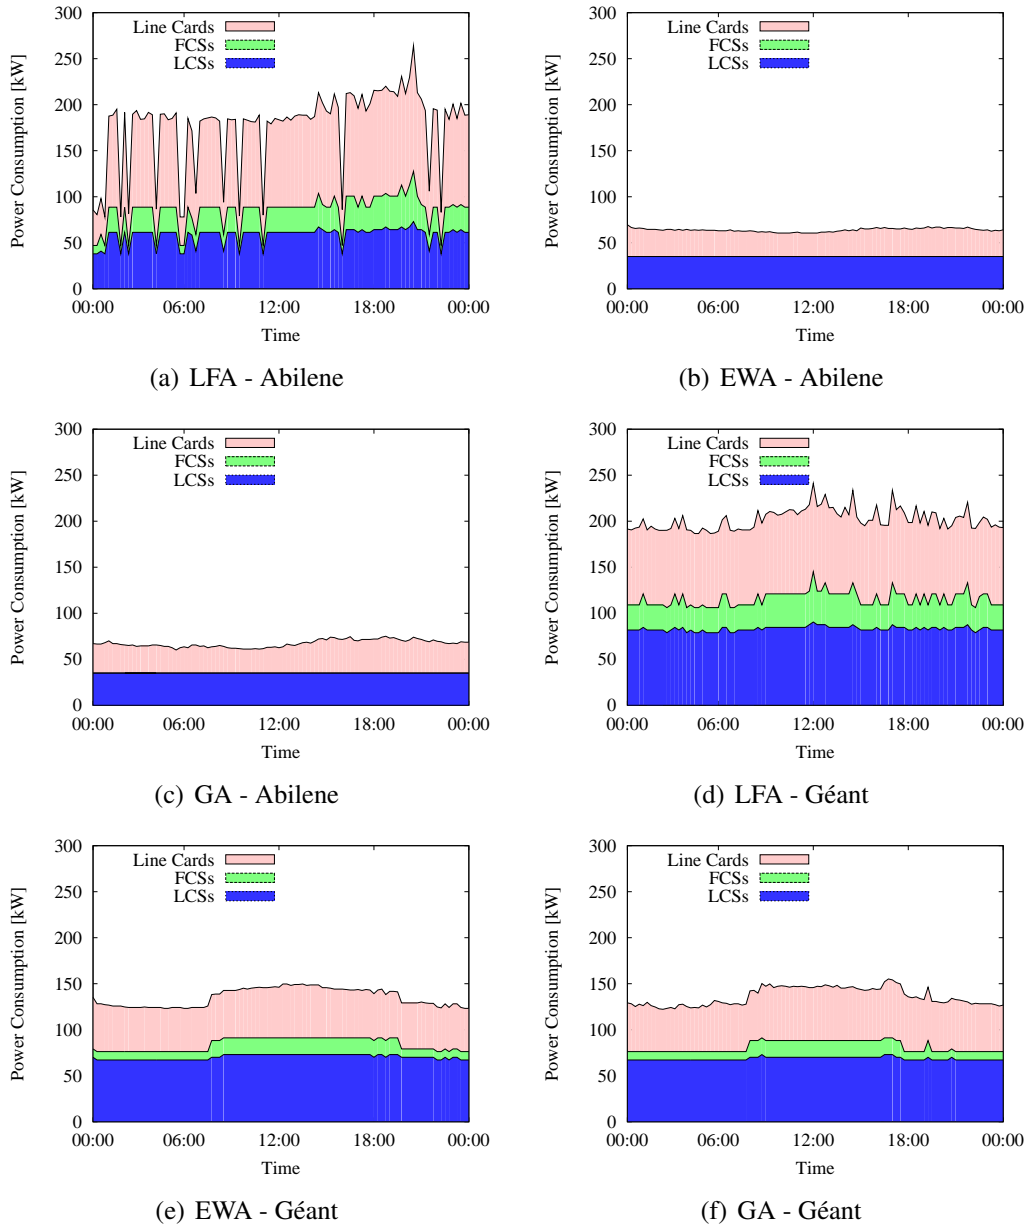


Figure 3.9: Power Breakdown for the Abilene and Géant networks.

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 LCs 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 Section 3.3.2), and significantly higher

Table 3.6: Yearly Monetary Cost [k€]

	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%)

number of installed LCs than LCSs.

### 3.5.5 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<sup>5</sup>. A cost of 0.0936 € per kWh is assumed, based on current statistics of electricity prices in Europe [32]. Table 3.6 reports the costs for the three algorithms. Although not being the main focus of this research activity, percentage savings with respect to an always on solution are reported in the parenthesis. Considering the Abilene network, an always on solution requires 280000 € per year. With LFA this amount can be almost halved, while with EWA and GA savings larger than 80% can be achieved, with 52000 € 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.

## 3.6 Discussion and Implementation Issues

Power savings and amount of reconfigured traffic depend on the traffic variation over time. We varied the load in Section 3.5.3, but the rate of traffic changes plays also an important role. The algorithms (especially EWA) are designed for typical day-night traffic variation in a backbone network. Traffic in this kind of networks is smooth due to aggregation of hundreds of end-to-end flows. Extraordinary and large peaks of traffic are unlikely, and caused by e.g., natural disasters or terrorist attacks, but also transmission of multiple

<sup>5</sup>Weekend days may further decrease the electricity consumption.

grid jobs. They may reduce the opportunity for putting devices into sleep mode, and in such cases, we expect lower power savings with a higher amount of reconfigured traffic. Besides, these situations may also result in network overload if the devices that were put into sleep mode cannot be waked up fast enough, or if the network is improperly dimensioned. Lastly, the algorithms may fail to find a suitable network configuration for supporting these peaks of traffic, which may impair the network availability. These issues and the exact domain of validity of the algorithms with respect to traffic dynamics have not been studied within this research activity.

In our case, we have evaluated our heuristics over two weekend days and two working days (for the Géant and Abilene networks), and one day for the Germany17 network. We have selected these days, because of their typical day-night traffic variation in a backbone network. We assume that in normal conditions, traffic in the network varies with a similar trend, and so the power saving and the amount of reconfigured traffic would be similar. Moreover, the adopted time granularity (one TM every 15 minutes) is sufficient to mimic the real traffic variation in the network.

Furthermore, there are different implementation issues that need to be taken into account when introducing the proposed heuristics into the networks. First of all, the feature of remote activation and deactivation of the LCs, and potentially also LCSs and FCSs has to be integrated into IP routers. Although commercial routers available today do not possess this functionality yet, their vendors work on its implementation [33]. In general, estimating the time that is needed to power up a LC in an IP-over-WDM network is not trivial. However, the 3 microseconds needed to wake up a 10 Gbps link in the emerging 802.3az standard [34] makes us believe that the wake-up process will be quick enough to follow the traffic variation.

Additionally, in the case of a centralized controller, it is necessary to consider the time needed to run the algorithm and the time to disclose the new network configuration to all the nodes. This can be done using a new signaling mechanism or an already existing one. For instance, the computed network configuration can be put inside Link State Advertisement (LSA) packets, assuming the existing MultiProtocol Label Switching (MPLS) as the control plane [35]. Alternatively, the algorithms can be run in a distributed manner at each node, which requires prior collection of the necessary global data.

Lastly, additional time is needed to reconfigure the network, first switching on new equipment, then making sure that all the traffic has been rerouted, and eventually switching off idle equipment. The make-before-break mechanism [36] can be used for this purpose.

## 3.7 Conclusion

We have targeted the reduction of power consumption in IP-over-WDM networks, by selectively turning off LCs during periods of low traffic. In particular, we have first formulated the Multi-Period Power-Aware Logical Topology Design problem (MP-PA-LTD) as an optimization problem. Then, we have proposed three different heuristics to solve the MP-PA-LTD, namely LFA, GA and EWA. The LFA targets the minimization of the number of logical links in the network, while both GA and EWA work on single lightpaths. Differently from the previous work, we take into account the amount of traffic which is reconfigured in the network.

We have evaluated our algorithms over a wide set of network scenarios taking as realistic parameters as possible. Our results indicate that switching off LCs is of great benefit to reduce the power consumption of the network. Moreover, both GA and EWA are able to wisely trade between power and reconfigured traffic, and this trade-off can be controlled by a proper setting of the input parameters. Then, we have shown that with our heuristics the traffic load impacts the power consumption, 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.

In conclusion, our results indicate that large savings are possible while keeping the reconfigured traffic low. This is true especially for the GA and EWA algorithms, which are designed to switch off single lightpaths rather than full logical links.



## Chapter 4

# Switching Technologies in Core and Metro Networks

The increasing of traffic in core networks, and consequently of the energy consumption, is motivating to search new solutions to improve the energy efficiency of the networks. A viable solution may be to use new switching paradigms that can reduce the consumed energy with respect to currently adopted switching techniques. In the following, we focus in particular on time-domain optical sub-wavelength switching technologies and we evaluate if they represent an effective solution for reducing the energy consumption.

Indeed, optical sub-wavelength switching technologies have been proposed to reduce the cost and the energy consumption, since they can exploit more efficiently the photonic resources [37]. These technologies can transparently switch traffic at a granularity finer than the wavelength, permitting to share the same interfaces to send/receive traffic to/from different destinations/sources.

This characteristic is particularly interesting, especially in networks dealing with several small flows such as backhaul networks. Indeed, it would permit, intuitively, to reduce the number of interfaces and, thus, the associated cost and energy consumption. However, in the literature [38], there is no clear consensus on the expected benefits of these technologies, probably due to the several different solutions that can be found.

In this chapter, we analyze three of these switching technologies and we compare them with legacy circuit switching technologies in order to understand in which conditions they are interesting and which features are the most advantageous.

The chapter is organized as follow. The circuit switching technologies are introduced in Sec. 4.1, while in Sec. 4.2 the selected optical sub-wavelength switching technologies are described. The network scenario and the adopted power consumption model are presented, respectively, in Sec. 4.3 and Sec. 4.4. The obtained results are discussed in Sec. 4.5 and conclusions are drawn in Sec. 4.6.

## 4.1 Circuit Switching Technologies

Nowadays, transport networks rely on circuit switching technologies. Different circuit switching paradigms have been introduced during the years. The first one has been opaque circuit switching which is based on optical to electrical to optical (O/E/O) conversions of all the traffic arriving at a node. Indeed, regardless the final destination of the traffic, each node processes electronically the traffic in order to forward or drop it. This process is very energy consuming and the large increase of the traffic volume has made the energy requirements even larger. Thus, the introduction of transparency in the network has been thought as a possible solution to overcome the electronic traffic processing, and also to reduce the energy consumption.

The term transparency refers to the direct optical transmission of data from source to destination using a logical optical channel, which is usually called lightpath. Thus, the traffic is switched at the optical layer at the intermediate nodes, and it is processed electronically just at the destination nodes.

This solution, however, if on the one hand eliminates the electronic traffic processing at transit nodes, on the other hand has the drawback of having a scarce aggregation capability. Indeed, the capacity of lightpaths is characterized by a coarse granularity, equal to the transmitters' capacity. Therefore, there is usually a large mismatch among the transmission capabilities between two nodes and the actual traffic requirements, resulting in a large underutilization of the resources.

A possible solution is to allow the grooming of different traffic demands in the same lightpath and let some demands span over consecutive lightpaths, following multi-hop paths, until the destinations are reached. In this case, it is required to perform O/E/O conversions and to electronically process the traffic demands following multi-hop paths, but this is just a relatively small fraction of the total traffic present in the network. This solution, called translucent circuit switching, can ameliorate the aggregation capability. Thus, it is possible to achieve a more effective use of the resources which allows to reduce the number of employed devices and consequently to reduce the energy consumption.

Thus, we want to understand if optical sub-wavelength switching technologies can further improve the energy efficiency of transport networks, thanks to better traffic aggregation capabilities and to not resort to electronic traffic processing along the entire transmission path.

## 4.2 Time-Domain Optical Sub-Wavelength Switching Technologies

We refer to time-domain optical sub-wavelength switching technologies as to the switching technologies that allow the nodes to exchange data, in the format of bursts of packets, sharing the same wavelength in the time-domain. In particular, we consider technologies



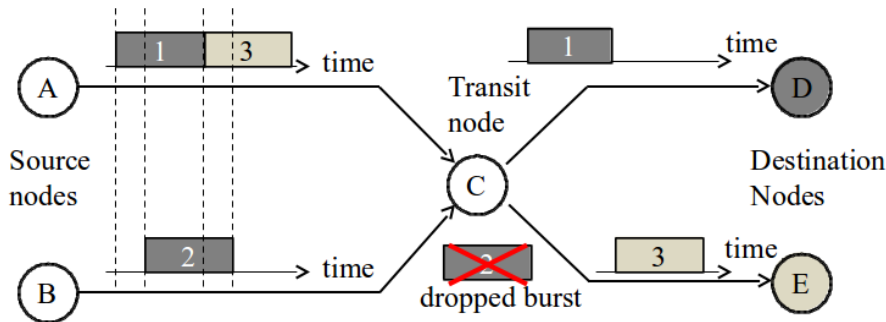


Figure 4.1: OBS functioning.

in which one or more wavelengths are uniquely assigned to a specific destination. Hence, the switching is performed at the source node by tuning the wavelength of the transmitter according to the destination (i.e., wavelength-routed technologies).

Thus, each node can use the same network interface to communicate with the other nodes, i.e., a node can use the same transmitter to forward traffic to different destination nodes and it can receive traffic from different nodes over the same receiver. In this way, a transmitter of a node and a receiver of another given node are not constrained to communicate only one with the other, but they can be freely matched with other network interfaces according to the current needs. This feature may permit to achieve a very efficient traffic aggregation and, together with the transparent transmission of data, it is likely to reduce the energy consumption with respect to circuit switching technologies.

Main issue of the sub-wavelength switching technologies is the way that the sharing of the wavelength is managed. Two main categories of these technologies are considered: lossy and loss-less. Lossy solutions do not guarantee the successful transmission of data since nodes can send bursts that can compete for the same wavelength at the same time. In this case, a contention happens and, depending on the contention resolution method, one or more bursts can be discarded.

In the other category, loss-less solutions adopt end to end reservation of the optical resources along the path, such that contentions and, thus losses, are not possible. The reservation is usually performed by scheduling the transmission of the bursts according to well defined schemes.

In the following sub-sections, we briefly describe the sub-wavelength switching technologies that we have considered. We refer to the previous works for a detailed description of each technology.

### 4.2.1 Optical Burst Switching

There are many different implementations of Optical Burst Switching (OBS). For our analysis, we have chosen a wavelength-routed and lossy version inspired from Label

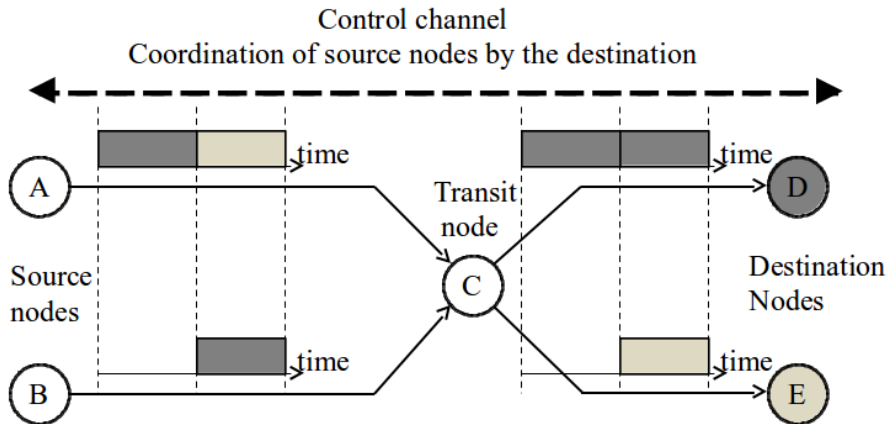


Figure 4.2: TWIN functioning.

OBS [39]. More precisely, a source node sends data by properly tuning the transmission wavelength as soon as a transmitter is freed. Thus, resources are not reserved in advance and there is no coordination between nodes, which may result in contentions and then losses of bursts as in the example of Fig. 4.1. In the case of contention, only one burst among those that request simultaneous access to the same resource is accepted, the others are discarded.

This solution, finally, needs some loss recovery process, which may be, for instance, burst retransmission. This requires increasing the transmitting capabilities of the nodes, resulting in an over-dimensioned network with respect to the real traffic requirements.

## 4.2.2 Time-domain Wavelength Interleaved Network

Time-domain Wavelength Interleaved Network (TWIN) [40] is a loss-less, wavelength-routed solution in which transmissions of bursts are scheduled according to time slots. The scheduling ensures that just one node at time can transmit on a slot for a given wavelength. In this implementation of TWIN, we consider that the scheduling is performed by the destination nodes.

The coordination among several sources and a destination is achieved using a request-grant scheme: the sources send their requests asking for transmission resources (i.e., time slots) to the destination using a separate control channel. The destination, considering all the received requests, allocates slots to each source and it informs them, through a grant message, which are the slots that they can use for transmission (see Fig. 4.2).

A disadvantage of this scheme is that, since destinations are not coordinated, a source can receive a grant for the same slot from several destinations. In this case, if the node is not equipped with enough transmitters, it cannot transmit to all the destinations. This issue is known as burst blocking [40] and it causes a degradation of the transmission

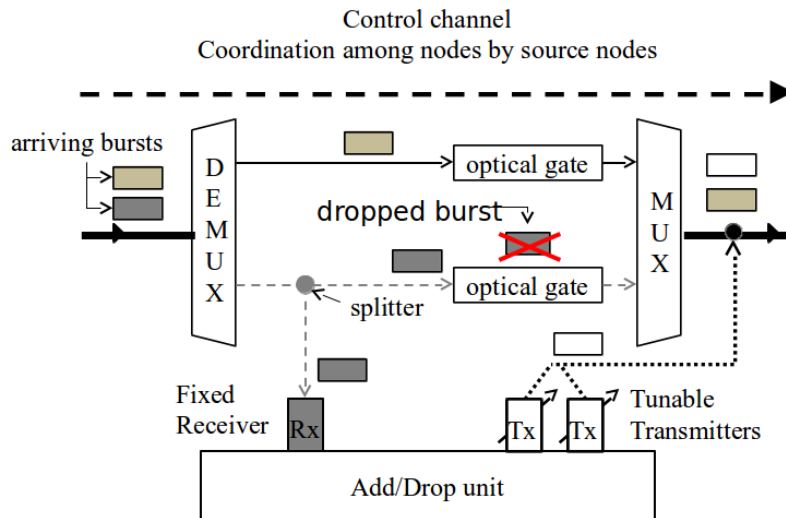


Figure 4.3: POADM simplified structure and functioning.

capability of the nodes.

### 4.2.3 Packet Optical Add Drop Multiplexer (POADM)

This switching technology has been conceived just for ring networks. It may be applied also to other physical topologies, but it is necessary to treat them, logically, as a ring topology.

The nodes of POADM, differently from the previous technologies, are not based on wavelength selective switches, but they use systems based on POADM structure [41] as illustrated in Fig. 4.3. Thanks to this structure, each node can receive bursts over a fixed set of wavelengths and it can transmit on any wavelength.

Similarly to TWIN, the time domain is slotted, but the transmissions to a given destination are scheduled using a different distributed reservation approach described in [42]. The reservation of slots for the transmission of bursts is performed by control messages on a separate wavelength channel according to scheduling cycles.

## 4.3 Network Scenario

We present the dimensioning results for a unidirectional ring network of 10 nodes and for a uniform traffic scenario. This traffic corresponds to a distributed traffic repartition envisioned in the case of Content Delivery Networks (CDNs).

The comparison among the technologies is based on the amount of transmitters (Tx) and receivers (Rx) which each node has to be equipped with in order to achieve a desired

flow per node. We define as flow per node the total amount of traffic, in Gbps, that a node can successfully receive from another node.

The dimensioning for the opaque and for the transparent switching technologies is retrieved analytically. Indeed, for a unidirectional ring network of  $N$  nodes, in the opaque case,  $F = N \cdot (N - 1)/2$  flows transit on each link. Defining  $\lambda$  Gbps as the traffic arrival rate of each flow and  $C$  Gbps as the capacity of Tx (and Rx), each node has then to be equipped with  $\lceil F \cdot \lambda / C \rceil$  Tx (and Rx).

In the transparent case, each node is sending directly the traffic to all the destinations. Therefore, the number of Tx (and Rx) required per node is equal to  $\lceil \lambda / C \rceil \cdot (N - 1)$ .

In the case of translucent circuit switching, it is necessary to perform the design of the network in order to choose which lightpaths must be established and how traffic demands are routed into the lightpaths. We utilize for that a meta-heuristic based on a genetic algorithm [43] which has been modified in order to minimize the total number of Tx and Rx in the network.

The number of Tx and Rx, for the sub-wavelength switching technologies, is determined using simulations. We model each sub-wavelength technologies in the OMNeT++ Network Simulation Framework. The following parameters have been used for the simulations: bursts of data have fixed size equal to  $b = 4 \cdot 10^4$  bits, at each node the bursts for a given destination arrive according to a Poisson process with arrival rate equal to  $\lambda/b$  and the capacity  $C$  of Tx and Rx is set to 10 Gbps. The time slots have a fixed duration equal to the duration of a burst plus a guard time equal to 5%, in order to take into account laser tuning time and synchronization accuracy issues.

In TWIN and POADM, since there are no losses, we dimension the network considering, at each source, a traffic arrival rate per destination  $\lambda$  equal to the desired flow per node. In the case of OBS, we choose to dimension the network by increasing the arrival traffic rate until the losses are compensated, in such a way that at destination nodes the flow per node is equal to the desired one. In this way, we emulate the retransmission and we take in account the additional resources required for it.

## 4.4 Power Consumption Model

In order to evaluate the energy efficiency of each solution, we propose the following power model for the Tx and Rx of circuit and sub-wavelength switching technologies.

We consider that the overall power consumption of a Tx and a Rx employed in circuit switching corresponds to the consumption of a full transponder (including FEC, client ports, framers,...) which consumes 40 Watt [44]. Within this value, the O/E/O conversion weighs for about 4 Watt, that is the power consumed by a typical XFP (100 Gigabit Small Form Factor Pluggable) [45]. We estimate that this power consumption can be attributed to the Tx, since the XFP consumption is mainly due to the laser. Thus, the power consumption of Tx and Rx employed in circuit switching is respectively of 22 W

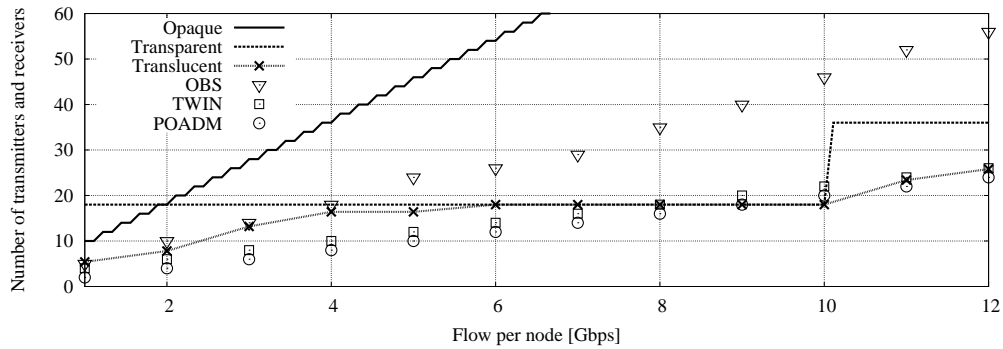


Figure 4.4: Number of transmitters and receivers per node vs. flow per node

and 18 W.

Focusing on sub-wavelength switching technologies, Rx require burst mode receivers which are consuming the same power of non burst mode receivers utilized by the circuit switching solutions [46].

Instead, the Tx of sub-wavelength switching solutions include a fast tunable laser, while circuit switching Tx employed slow tunable lasers. Both types of lasers are similar except that fast ones require adequate electronics to support the faster tunability. Comparing the power consumption of some fast [47, 48] and slow [49] tunable lasers, that are commercially available, we have that both present a similar power consumption comprised between 3 and 4 Watt. However, the fast electronic control circuitry may increase of several Watts the power consumption of the Tx employing fast tunable lasers. We evaluate that this increase is equivalent to at most double the laser power consumption, thus it results that Tx employed in sub-wavelength switching technologies consumes 26 W, at most 4 Watt more than conventional Tx.

## 4.5 Results and Discussion

The total number of Tx and Rx required by a node to achieve a given flow per node (i.e., the traffic that a node is receiving by another single node) is depicted in Fig. 4.4.

The figure shows that opaque circuit switching requires limited resources just for very low traffic loads, while, as soon the traffic load grows, the required resources steeply increases. Instead, transparent switching achieves interesting results just when the flow per node is close to the Tx capacity. Indeed, the drawback of transparent switching is that it needs a significant initial number of Tx and Rx, since data are transmitted directly among sources and destinations.

As expected the translucent switching solution is always performing better with respect to the opaque and transparent cases. Indeed, it is able to choose, depending on

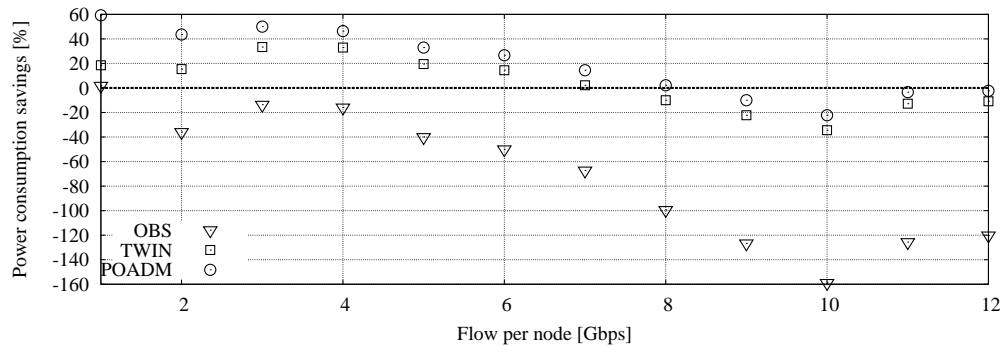


Figure 4.5: Power consumption savings of sub-wavelength technologies with respect to translucent circuit switching

the traffic load, which is the best trade-off between opaque switching and direct optical transmission.

The lossy OBS performs well only at low loads. Indeed at high loads, the over-dimensioning required to recover from the losses has a real detrimental effect on the dimensioning results. Thus, the absence of coordination for the transmission seems to have no particular advantages, apart that no synchronization is required.

On the other side, TWIN and POADM are performing better than the translucent circuit switching at low loads, despite of the guard time between bursts and the distributed nature of their scheduling, while in the high load case, TWIN and POADM perform very close to the translucent solution. This can be explained by the fact that burst blocking in TWIN and non-optimal resources reservation in POADM result in waste of transmission capacity. A possible solution to these issues is to centralize the scheduling, but at the cost of complexity for the nodes (tighter control and synchronization) and of additional latency for the data (scheduling and signaling delays and optimization computation time).

The percentage of power consumption savings, due only to interfaces, of sub-wavelength switching technologies with respect to translucent circuit switching paradigm is plotted in Fig. 4.5. As expected, OBS can not achieve any savings and its power consumption steadily increases with the traffic load. POADM and TWIN can offer up to 60% and 35% power savings respectively, when traffic load is low. As the traffic increases, savings are diminished until translucent switching becomes more energy convenient. Indeed, even if the number of Tx and Rx is similar, the slightly larger power consumption of fast tunable lasers is making sub-wavelengths technologies not convenient. This result however does not take into account the nodes consumption due to each particular architecture.

## 4.6 Conclusions

We have analyzed two different categories of switching technologies: circuit and time-domain optical sub-wavelength switching technologies. We have shown that, in a low and uniformly distributed traffic scenario, optical sub-wavelength switching technologies can improve the energy efficiency of networks with respect to legacy circuit switching.

In particular, we quantify that sub-wavelength switching technologies can reduce the power consumption of networks on average of about 30%, in low traffic case, respect to circuit switching. This means that these technologies could be a promising solution to reduce the power consumption of core and metro networks.





## Chapter 5

# Energy Efficiency in Access and Aggregation Networks: Status and Perspectives

Access and aggregation networks consume nowadays a large share of the energy consumed in wired telecommunication networks. Some actions to ameliorate the energy efficiency of these network segments have thus to be undertaken in order to ensure the scalability of future networks.

In this work, thus, we want (i) to characterize the current energy situation and (ii) to evaluate the effectiveness of applying energy efficient policies in these network segments, considering realistic scenarios. For this purpose, we take advantage from a wide dataset of measurements collected in the network of FASTWEB, a national-wide operator in Italy and partner of the TREND project.

This dataset comprises energy consumption, traffic and other important measures of its access and aggregation networks. Thus, using these actual data, we characterize the energy consumption of these network segments in order to understand how and where the energy efficiency can be improved. Next, we perform a “what-if” analysis in which we identify possible energy efficient policies that can be implemented if new technologies become available or if different management procedures are adopted. Moreover, the energy savings achievable by these policies have been precisely evaluated using the real data collected by FASTWEB.

In the context of access and aggregation networks, we focus our attention on the Point of Presence (PoP), large network nodes that act as both aggregation and traffic switching points, and on ADSL modems, located at the users and at the Digital Subscriber Line Access Multiplexer (DSLAM). This choice is motivated by the fact that the collected data show that the energy consumption of PoPs weights for the 26% on the total energy consumed by FASTWEB during one year. The quota of energy consumed by users terminals in access networks is also significant: it has been stated that they represent more than 65%

of the overall energy consumed in access networks [12].

Notice that the PoPs and the ADSL modems operate in substantial different traffic scenarios, indeed the PoPs aggregate the traffic of several users, while the traffic managed by the modems is related to single users. Thus, we have to take into account this when we conceive possible policies to improve the energy efficiency. In particular, in the PoPs, the traffic variations are smooth and we expect to have a daily traffic behavior similar for the different PoPs. Instead, in the case of ADSL modems, the traffic depends on the activity of just one user. Thus, the traffic variation are very bursty, since we switch from a state in which traffic is present (i.e., the user is active) to another in which there is no traffic (i.e., the user is not active). Moreover, each line has a different traffic profile, according to the habits of the user.

In the following, the chapter is organized: in Sec. 5.1 we briefly describe the dataset and the available metrics. In Sec. 5.2 we perform the energy profiling of the PoPs, evaluating possible correlations among energy consumption and the other metrics, and furthermore we introduce the energy breakdown of the PoPs. We analyze the traffic dynamics of the PoPs and we evaluate the typical activity of ADSL user in Sec. 5.3. The energy saving strategies are introduced in Sec. 5.4, while the achievable energy savings are estimated in Sec. 5.5. Lastly, conclusions are drawn in Sec. 5.6.

## 5.1 Dataset

The available dataset comprises an extensive set of informations regarding the main characteristics of the PoPs as well as measurements of the energy consumed by the PoP, the outside temperature and the traffic activity of both PoPs and single ADSL lines.

### 5.1.1 PoPs' data

The informations characterizing the PoPs consist in the number of users associated to it, the number of networking devices hosted by the PoP, which access technology is employed by the users and the type of the air conditioning system. A brief overview of these characteristics is given in Sec. 5.2.1.

Since FASTWEB has several PoPs, in the order of some tens, we select just a small subset of eight PoPs for our analysis. We have chosen the PoPs in such a way that they present different characteristics for what concerns the dimension, the geographical location, the access technologies, the air conditioning system. In our analysis, indeed, we want to consider an heterogeneous subset in order to understand which are the factors that impact on the energy consumption.

The energy consumption and the outside temperature measurements have been collected thanks to a management system that FASTWEB has implemented recently. The objective of this system is to monitor the energy consumption of company facilities in

order to check if energy inefficiencies, due to failures or to wrong system configurations, are present.

We collected measurements of the energy consumed and of the temperature in the period between mid September 2010 and mid September 2011. In this way, the set of measurements comprises all possible seasons.

In details, the system provides the:

- **Total energy consumed by a PoP**, it comprises the energy consumed by all the devices, including network equipments and air conditioning system devices, that are installed in a PoP. The energy is expressed in kWh and the measures have a granularity of the quarter of hour. Denote by  $\Delta$  the time granularity,  $\Delta = 15$  min.  $E_X(n)$  is the energy consumed by PoP  $X$  during the  $n$ -th time window since the beginning of the measurement campaign. To avoid outliers, we restrict the set of measurements to the subset of  $E(n)$  values within the 1-st and 99-th percentile of  $E_X(n)$  distribution. We denote this set with  $\mathcal{E}_X$ .

- **Air temperature** at the external of the PoP. It is measured with a granularity of an hour, but for ease of notation we denote it by  $T_X(n)$  referring to time windows of length  $\Delta$ , defined as above. In this way, we associate to each energy consumption sample the sample of the temperature collected in the same hour.

Concerning the measurements of the traffic processed by the PoP, we consider the period starting from January 2011 until the end of March 2011. We restrict our analysis on traffic to this period, since the number of users at each PoP can be considered stable, while it typically changes during other periods of the year when customers are attracted by special offers. Each traffic sample consists in the:

- **Normalized bit rate** of traffic processed by a PoP. The system exposes  $B_X(n)$ , the total amount of traffic processed by PoP  $X$  in  $\Delta$  time, which corresponds to the sum of all the data that are going from the users associated to PoP  $X$  to the network and vice-versa. Thus, the average bit rate in time window  $n$  is given by  $B_X(n)/\Delta$ . Finally, we normalize the bit rate to an arbitrary value that we choose to be equal to 0.8 the maximum observed bit rate. That is,

$$\rho_X(n) = 0.8 \frac{B_X(n)/\Delta}{\max_i(B_X(i)/\Delta)} \quad (5.1)$$

This is equivalent to assuming that the PoP capacity has been dimensioned so that the traffic load does not exceed 0.8.

### 5.1.2 ADSL lines' data

We collect traffic data of about 9000 ADSL residential customers, associated to the same PoP, for 24 hours, starting from 31 May 2011 at 11pm. At the end, 16 millions of samples have been collected, corresponding to an aggregate of 7.6 TB of bitwise volume.

The traffic measurements of each single ADSL line consist in:

- **Number of bytes** sent and received by the user in a time interval of 10 seconds. These samples have been collected using the traffic monitoring tool Tstat [50,51]. Starting from the measured bytes, we compute the line rate, expressed in b/s, as  $bytes \cdot 8/10$ .

## 5.2 Energy consumption characterization of the Points of Presence

Table 5.1: Summary of PoPs characteristics

ID	Users [k]	Devices	FTTH [%]	Free cooling [%]	$r_{Temp}$	$r_{Load}$
<i>A</i>	3.8	6	100	100	0.9	0.24
<i>B</i>	52.2	31	33	0	0.47	0.11
<i>C</i>	22.8	35	0.2	12	0.81	0.09
<i>D</i>	16.3	19	86	0	0.8	0.08
<i>E</i>	13.5	34	92	0	0.75	0.28
<i>F</i>	62.3	72	11	40	0.75	0.35
<i>G</i>	61.9	44	6	100	0.85	0.36
<i>H</i>	22.5	32	0.2	4	0.75	0.13

In this section, after a brief overview of the main characteristics of the PoPs, we characterize their energy consumption, evaluating the correlation between the energy consumption and the other metrics. Then, we analyze the energy consumption breakdown of each PoP, in order to determine the energy efficiency of the networking devices and of the air conditioning system.

### 5.2.1 Overview of the main characteristics of the PoPs

The main characteristics of the PoPs are summarized in Table 5.1. In particular, we report the characteristics of just the eight PoPs that we selected. Hereafter, these PoPs are identified with a capital letter ranging from *A* to *H*, as it is possible to notice in the first column of the table.

PoPs have different dimension depending mainly on the area that they cover. We quantify their size using as reference the number of users associated to each PoP and the number of hosted networking devices. These informations are respectively reported in the second and third column of Table 5.1. The number of users ranges from a few thousands for PoP *A* to more than 60 thousands for *F* or *G*. Similarly, the number of devices ranges from a few units of PoP *A* to more than 70 of PoP *F*.

The access technologies that users can be provided with are ADSL and Fiber-To-The-Home Ethernet-based technologies. Depending on the infrastructure that has been

deployed in the surroundings of the PoPs, both or just one between ADSL or FTTH systems can be present in a PoP. The percentage of users in a PoP that are using FTTH access technology is reported in the fourth column. Note that typically the PoPs have a predominant access technology.

FASTWEB has deployed, inside the PoPs, air conditioners belonging mainly to two different categories. The first category includes air conditioners which cool the air only using heat pump systems, usually based on compressors. The other category comprises air conditioners which are provided also with the free cooling option. In these devices, when the external temperature is sufficiently low, there is the possibility to cool the inside of the PoP using the outside cold air which allows to cool down consuming less energy. In each PoP, there may be present air conditioners of only one category or they can be mixed. The percentage of cooling capacity of a PoP, obtained using air conditioners with the free cooling option, are reported in the fifth column of the table.

The last two columns of Table 5.1 report the Pearson correlation coefficients of the energy consumed by a PoP in relation, respectively, with the external temperature and with the traffic processed. In the following, we describe in details the obtained results.

## 5.2.2 Energy consumption correlation with external temperature and traffic load

The analysis is based on the computation of the sample Pearson correlation coefficients which indicate the correlation between the energy consumption and another metric. In particular, we determine if there is any correlation between the energy consumption and the outside temperature, and between the energy consumption and the traffic load.


The sample Pearson correlation coefficients, denoted by  $r$ , have been computed using Eq. 5.2, where  $x_i$  and  $y_i$  represent, for each metric, one sample out of the dataset of dimension  $N$  and where  $\bar{x}$  and  $\bar{y}$  are the average value.

$$r = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2} \cdot \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}} \quad (5.2)$$


If the coefficients  $r$  assume value equal or near to 0, it means that the two metrics have none or little correlation, instead if values are close to 1, the two metrics are linearly correlated.

### Energy consumption correlation with external temperature

the sample Pearson correlation coefficients of the energy consumption with the external temperature are reported in the sixth column of Table 5.1. It is possible to notice that there is a strong correlation for most of the PoPs and, in particular, the correlation is stronger in



(a) PoP A



(b) PoP B

Figure 5.1: Energy consumption and external temperature versus time.

the PoPs where the percentage of air conditioners with the free cooling option available is larger.

We take as exemplary PoPs, the PoP A and the PoP B, which are respectively the PoP with the highest and the lowest correlation coefficient, and they are also the PoP with full and with none free cooling capability. Figs. 5.1(a) and 5.1(b) report the energy and temperature of the whole measurement campaign for both PoPs. Each plot is composed of two parts: the top part shows energy consumption,  $E_X(n)$ , the bottom part the temperature,  $T_X(n)$ . Looking at Fig. 5.1(a), it can be noticed that in PoP A the trend of the energy consumed follows quite closely the temperature trend. In particular, during the coolest days of the year, the energy consumed reaches its minimum values; when spikes in the temperature are present, the energy consumed is larger. Energy and temperature




Figure 5.2: Normalized energy consumption vs. external temperature for PoP A.




Figure 5.3: Energy-temperature correlation coefficient vs. free cooling capability.

follow the same seasonal behavior, with low values from October to March<sup>1</sup>. On the contrary, Fig. 5.1(b) shows that, for PoP B, the consumed energy barely follow the yearly temperature trend.

The correlation between energy consumption and temperature can be perceived also in Fig. 5.2 which report the normalized energy versus the external temperature for PoP A. The normalized energy,  $\hat{E}_X(n)$ , is used in order to avoid any influence due to the different

<sup>1</sup>The sudden jump at beginning of April in Fig. 5.1(a) is due to a change in the configuration of the cooling system.




Figure 5.4: Normalized energy consumption versus traffic load.

sizes of the PoPs. It is defined as,

$$\hat{E}_X(n) = \frac{E_X(n) - \min(\mathcal{E}_X)}{\max(\mathcal{E}_X) - \min(\mathcal{E}_X)} \quad (5.3)$$

for the samples  $E_X(n)$  that belong to  $\mathcal{E}_X$ .

The impact of the air conditioning system type over the correlation can be evaluated in Fig. 5.3, which depicts the correlation coefficient  $r_{Temp}$  of each PoP versus the percentage of free cooling air conditioners capacity. PoPs *A* and *G*, that use only free cooling conditioners, have the highest correlation coefficients, while PoP *B*, which do not use at all the free cooling option, has the lowest one, even if it is still quite high. Other PoPs, which have none or little free cooling capability, present a larger correlation with respect to PoP *B*. Thus, the free cooling capability has a relevant impact over the correlation between energy and temperature, but there are also several other factors, such as the exposure and the physical location or how the building hosting the PoP is built, that have an impact on the efficiency of the air conditioning systems.

### Energy consumption correlation with traffic load

the correlation between the energy consumed and the traffic load has been retrieved similarly to the previous case. The values of the correlation coefficients  $r_{Load}$ , reported in the last column of Table 5.1, are much smaller than one, meaning that there is little correlation between these two metrics.

This fact is confirmed by Fig. 5.4, where the normalized energy is plotted against the normalized traffic for PoP *G* which is the PoP that presents the largest correlation coefficient among the selected PoPs.



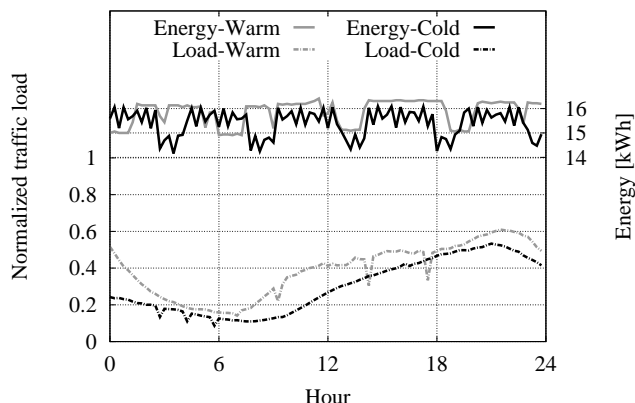


Figure 5.5: Energy consumption and normalized traffic profiles of PoP *B* during one day.

We can retrieve the same conclusions looking at the daily variation of consumed energy and traffic load of PoP *B* that are plotted in Fig. 5.5. We report the daily variations for two days, one day during the cold season (black lines) and one in the warm season (gray lines). Left y-axis reports the normalized traffic (dashed lines), while right y-axis reports the energy consumption (solid lines). While the traffic follows always a day-night pattern, with just some slight differences between the two days, the energy consumption profiles present a completely different behavior with respect to the traffic.

Moreover, some interesting observations, about the correlation of energy and external temperature, can be deduced. The energy consumption presents an intra-day periodic behavior that can be explained by the switching on and off of the air conditioner compressors according to a well defined scheme. In the warm day, furthermore, an additional effort is required to cool down the PoP. Indeed, it is possible to notice that the number of peaks are smaller, since the conditioners are switched off less frequently, and there is also a small increase in the maximum value of the consumed energy.

### Energy consumption correlation with other indexes

we have also studied the correlation between the energy consumption of the PoPs and other parameters. In particular, we have considered: (i) the size of the PoP in terms of number of users, and (ii) the predominance of ADSL or FTTH technology.

Fig. 5.6 shows the average energy consumption per user. Measurements do not allow to conclude that there is a correlation between energy consumption and size of the PoP. For example, in both PoP *A* and *B*, the average consumption per user is about 0.5 Wh despite PoP *A* accounts for only one tenth of the users served by *B*. Notice that the high values of PoP *E* and *F* are due to the presence in those PoPs of WDM equipments.

Finally, we do not notice any trend that can be associated to the access technology. Results are not reported here for the sake of brevity.




Figure 5.6: Energy consumption per user versus number of users.

### 5.2.3 Energy breakdown analysis

Table 5.2: Energy breakdown of the PoPs

ID	$\min(\mathcal{E})$ [kWh]	$\max(\mathcal{E})$ [kWh]	$PUE$
A	0.7	3.7	5.29
B	14.28	16.75	1.17
C	10.34	14.56	1.41
D	6.18	7.53	1.22
E	13.63	16.84	1.24
F	41.4	59.42	1.43
G	21.44	29.13	1.36
H	12.41	17	1.37

The energy management system implemented in the PoPs, as already stated in Sec. 5.1, can just report the measurements of the total consumed energy in a PoP. It is not possible to differentiate the energy that is consumed by the air conditioners or by the networking devices, considering these as the two main contributors to the consumption of energy.

Anyway, it may be reasonable to assume that the minimum energy measured, considering a long time period, could refer to a moment of time in which all the air conditioners were idle and not consuming energy. This is a rough estimate, but it is important to understand how consistent is the part of energy that can be attributed to the networking equipments, since it is on this part that the energy saving strategies can mostly act. Indeed, if the consumption of the air conditioning system counts for a large part of the overall consumption, it may not be worth to implement energy saving strategies on the

networking devices, since the benefits will be negligible. On the contrary, if the energy consumption is mostly due to the networking devices, it may be convenient to search for new energy-aware strategies which may bring significant savings.

We thus evaluate the energy efficiency of the PoPs considering the Power Usage Effectiveness (PUE). This metric is computed as the ratio between the power consumed by the overall facility and the power consumed just by the IT equipment. In our case, we assume that these values are represented respectively by the maximum and the minimum value of consumed energy which are present in our dataset. In particular, for a PoP  $X$  we select  $\max(\mathcal{E}_X)$  and  $\min(\mathcal{E}_X)$ , from the set  $\mathcal{E}_X$ , that is the set of measurements within the 1-st and the 99th percentile of the  $E_X(n)$  distribution, in order to avoid outliers.

It is possible to notice in Table 5.2 that the PUE ranges from about 1.2 to 1.4 for most of the PoPs. Thus, the energy consumed by the air conditioning system and by all the non IT devices weights in a limited way over the total energy consumption.

Only PoP  $A$  has a very large PUE, meaning that the air conditioners have a drastic impact on the energy consumed by this PoP. This result can be explained by the fact that PoP  $A$  is a small PoP equipped with just six networking devices, as reported in Table 5.1. Thus, the air conditioners, that are required to cool down the facility, are dominating the energy consumption of the PoP since the IT energy consumption is very limited.

#### 5.2.4 Wrap-up

The energy profiling of the PoPs leads to the conclusion that the PoPs have an energy consumption that is correlated to the external temperature and, depending on the air condition system and other factors, this correlation can be weaker or stronger from a PoP to another one. Instead, none of the PoPs present an energy consumption correlated to the traffic processed.

These results motivate the effort to search for strategies that can make the consumed energy proportional to the traffic that a PoP manages. Furthermore, the energy consumption breakdown indicates that the networking devices are mainly responsible for the energy consumed in the PoPs.

### 5.3 Traffic characterization of Points of Presence and ADSL lines

The strategies for introducing energy-to-traffic proportionality can be effective depending on how the traffic changes over time. Indeed, if a PoP present very little traffic dynamics, i.e., it presents a small difference between the minimum and the maximum amount of processed traffic, the capability to adapt the energy consumption to the actual traffic processed is useless. Instead, if the PoP presents a significant variability in time of the




Figure 5.7: Typical daily pattern at a PoP along one week.

amount of traffic managed, it may be worth to have the PoP able to manage different traffic capacities with related energy consumption levels.

In the following, we perform a comprehensive analysis of the traffic processed by the PoPs, in particular to characterize the daily and weekly dynamics. Indeed, since a PoP is aggregating the traffic of a large number of users, it is expected that the traffic follows a daily profile without having abrupt variations.

Similarly, we profile the traffic load of ADSL lines in order to determine which energy efficient policies may be suitable for that context. In this case, since the traffic load depends on the activity of a single user, we expect that the traffic variations for each ADSL line are sudden and the traffic profile can vary significantly from a line to another one. For this reason, we perform the analysis focusing on a minutes time scale.

### 5.3.1 Macroscopic traffic analysis of the PoPs

Let us start by focusing on the traffic variation of the PoPs during a one-week long period. Fig. 5.7 reports measurements for PoP  $B$ , being the patterns for the other PoPs very similar. As can be seen, traffic follows a day/night periodicity, with off-peak traffic that is about one fourth of peak traffic. Only small differences can be noticed in different days.

To better quantify how traffic varies during the day, we operate as follows. Let  $\mathcal{R}_X(i)$  be the set of PoP  $X$  traffic measurements that correspond to day  $i$ . By sorting the elements in  $\mathcal{R}_X(i)$  by increasing values, we obtain the traffic distribution for day  $i$ . Fig. 5.8 shows the normalized traffic distribution for PoP  $A$ , which is representative of other PoPs. Measurements cover the 7 days of a week; Saturday and Sunday are highlighted. Notice that, for all the considered days, small differences during weekends can be observed.

Fig. 5.9 shows the normalized traffic distribution, for all PoPs, evaluated considering




Figure 5.8: Daily traffic distribution for a week of PoP A.

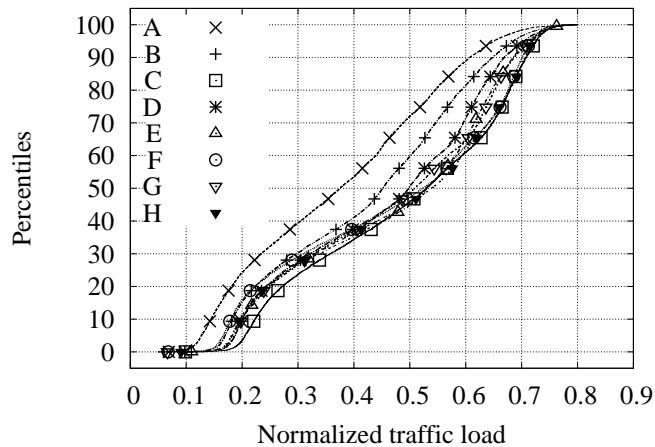


Figure 5.9: Traffic distribution of every PoP.

the whole 3-months long traffic dataset. It confirms that the distributions follow the same behavior for all the PoPs and also considering a longer time period. This means that the overall behavior of an aggregate of users is variable in a day in a quite stable and easily predictable way, regardless the location, size and kind of PoP.

Finally, the variability of traffic have been evaluated numerically computing the ratio between the peak and off-peak traffic values. Since the traffic is a noisy measure, and the actual value of the peak may depend on different local phenomenon, we compute for each day the ratio among the 80-th and 20-th percentile of the traffic measurements of the day. We retrieve that the daily excursion of traffic is very large, with typically 80% of peak hour traffic three times larger than the 20% of off-peak value.




Figure 5.10: Examples of the evolution of the line rate for three different users.

### 5.3.2 Microscopic traffic analysis of ADSL modems

Some examples of classic users' activity are reported in Fig. 5.10. In this figure, it is reported, using the logarithmic scale, the total line rate, obtained aggregating both download and upload traffic. These examples, while not being the only possible activity patterns, show that each user behaves in his own peculiar way.

In the figure there are shown three different usage patterns: *user1* is always connected constantly exchanging a high volume; *user2* has two activity periods during the morning and the afternoon; *user3* generates several traffic events, in an almost periodic way, characterized by very low volumes.

The traffic profile of *user3* corresponds to the case in which traffic is not generated by the user, but it is automatically generated by some applications or by the modem itself for control purposes. In the following, we use the term background traffic when we refer to this type of traffic. Note that also *user2* presents a traffic with similar characteristics when it is not generating significant volumes of traffic.

In the time periods in which only background traffic is present, users are not active and thus modems can be enforced to enter a sleep mode state, in which ADSL modems are switched off or put in a low power state in order to save energy. However, it is necessary to assume that the transmission of background traffic is not compulsory in those time instants and that it could be postponed until the user becomes again active.

The implementation of sleep mode policies in the ADSL modems can thus be a good solution to improve the energy efficiency in access networks, but their effectiveness depends on which is the typical user behavior. Indeed, if most of the users has a traffic profile similar to the one of *user1*, sleep mode policies are useless. In the case that the users' behavior is more similar to *user2* and *user3*, there is room to achieve significant

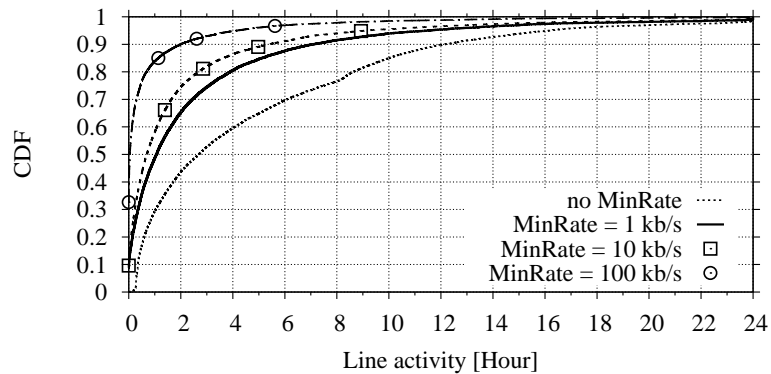


Figure 5.11: CDF of the total activity of the monitored lines during the day.




Figure 5.12: CDF of the utilization of the monitored lines during the day.

energy savings.

In order to understand which is the typical user behavior, we proceed evaluating which is the distribution of the activity periods of the users. We identify as activity periods the time periods for which the line rate of an ADSL line is above a certain minimum bit rate, called *MinRate*. In this way, we consider that all the traffic events for which the line rate is below *MinRate* are not to be associated to the user, but they represent background traffic.

In Fig. 5.11, it is reported the Cumulative Distribution Function (CDF) of the total line activity using different patterns for different values of *MinRate*. We can observe that, when no threshold is applied, only 50% of the lines is active for more than 2.5 hours and only 10% for more than 12 hours. The average line utilization is very limited. In fact, 50% of the lines are above 1 kb/s only for 1 hour, while this value drops to nearly 0 for 100 kb/s.

This phenomenon is even more clear if we consider the line utilization, defined as the ratio between the number of samples with bit rate higher than *MinRate* and the total samples of each line. Analyzing the CDF of the line utilization shown in Fig. 5.12, we

can see how only 30% of the lines reach 100 kb/s for at least 10 seconds during the day while samples above 100 kb/s account for only 25% of the samples of the line. Moreover, the distributions for 1 kb/s and 10 kb/s are very smooth suggesting that it is very common to have samples with very low bit rate.

### 5.3.3 Wrap-up

The traffic analysis of the PoPs shows that each PoP has a traffic profile which presents high dynamics. The traffic load in a PoP follows a regular day and night pattern with a relevant variation between the peak and the off-peak load. This means that, at the PoPs, it is convenient to introduce configuration strategies or new technologies which make the energy consumption proportional to the traffic that a PoP manages.

From the traffic analysis of the ADSL lines, instead, it results that most of users is inactive for long time periods, during the which only background traffic is generated. Thus, the implementation of sleep mode policies can be a good solution, at least for users that do not let P2P applications running for long periods of time.

## 5.4 Energy efficient policies

In this section, we describe some energy-wise strategies that can ameliorate the energy efficiency in access and aggregation networks. In particular, we consider that the energy reduction may be achieved (i) introducing *energy proportionality* at the PoPs, (ii) implementing *sleep mode policies* in the ADSL modems at the users' and DSLAM side.

The aim of this work is to evaluate which are the benefits that these strategies may bring. We do not focus on the problems related to their technical implementation. We just want to understand if using these strategies it is possible to obtain substantial energy saving that can motivate the choice to invest in this direction, as several manufacturers have already done.

### 5.4.1 Energy proportionality in the PoPs

The energy proportionality refers to the capability of enabling a networking device or a set of them to adapt the consumed energy to the actual traffic load. The analysis of the traffic dynamics, reported in Sec. 5.3, shows indeed that the traffic managed by the PoPs is varying in time, but the consumed energy does not change in relation to these traffic variations, as shown in Sec. 5.2.

The idea of energy proportionality is thus very appealing, since current devices consume energy not for the work they actually do, and the used bandwidth, but for the deployed capacity, as shown in [52]. Indeed, it is useless to have available a large capacity,






Figure 5.13: Linear and DVFS energy-to-traffic proportional strategies

which requires a large energy consumption, while the actual needs are limited to a small fraction of this capacity.

Several research projects are currently investigating how to achieve energy proportionality by implementing energy saving capabilities on the individual devices or the system architecture [8–11]. Moreover, since PoPs can be in a certain way considered as a special type of data centers, some research studies already analyze the energy consumption of data centers and they propose also strategies to increase their energy efficiency [53–56]. Nevertheless, no studies have proposed and evaluated solutions specifically for the PoPs.

In this work, we consider two different approaches to tackle the issue of improving the energy efficiency: (i) the availability of a new generation of networking devices which employ *energy to traffic proportional technologies*, (ii) a modular organization of the devices in the PoPs according to the *resource consolidation* practice [26, 56, 57].

Both these approaches can act on the consumption of networking devices only and not on the air conditioning systems. However, since the emitted heat depends on the energy consumption of the devices, it can be expected that some saving can be achieved also by the conditioning system.

### Energy to traffic proportional technologies

we assume that networking devices present an energy consumption which is fully or partially proportional to the traffic load.

The consumed energy  $e(\rho)$  is the sum of a fixed energy consumption,  $E_{static}$ , independent from the traffic and of a dynamic one,  $E_{dynamic}(\rho)$ , that is proportional to the traffic load  $\rho$ . Thus,

$$e(\rho) = E_{static} + E_{dynamic}(\rho) \quad (5.4)$$

In particular, we consider that the dynamic component of the energy is *linearly* or *cubically* proportional to the traffic processed. The latter one represents an approximation of implementing Dynamic Voltage and Frequency Scaling (DVFS) techniques in the devices. In Fig. 5.13, they are depicted exemplary behaviors of both cases.

In the following, we assume to use normalized values for the traffic load  $\rho$  and for the energy consumption values. The normalized traffic of the PoPs is computed according to Eq. 5.1. Moreover, we assume that the normalized energy consumption of the current technology (no proportionality) is equal to 1, and that is also the consumption at maximum traffic load for any energy proportional strategy, thus  $e(1) = 1$ .

Thus, in the case of *linear* proportional energy consumption, the actual consumed energy can be computed as:

$$e(\rho) = E_{static} + \alpha \cdot \rho \quad (5.5)$$

where  $\alpha \cdot \rho$  represents the variable component of consumed energy,  $\alpha$  being the proportionality factor.

Since we consider the normalized energy  $e(1) = 1$ , the fixed energy cost is  $E_{static} = 1 - \alpha$ . The case  $\alpha = 0$  corresponds to current technology, whose consumption is constant, regardless the load, and thus the normalized consumption is 1.

Similarly, when  $E_{dynamic}$  is *cubic* proportional to the traffic load, we have that:

$$e(\rho) = E_{static} + \alpha \cdot \rho^3 \quad (5.6)$$

Finally, according to the previous assumptions, the total energy savings  $S$ , that can be achieved, can be computed as:

$$S = 1 - \frac{1}{N} \sum_n e(\rho(n)) \quad (5.7)$$

where  $e(\rho)$  is the total energy consumed when processing the normalized traffic load equal  $\rho$ , with  $0 \leq \rho \leq 1$  and  $N$  is the total number of  $n$ , that are the sampled time periods.

### Resource consolidation

in this approach, the networking devices in a PoP are deployed with a modular organization and, during low traffic demand periods, a fraction of them can be kept active, while the remaining can be temporarily put into low consuming sleeping modes.

In this way, the deployed capacity (and the consumed energy) is not constant but adapts to the traffic needs. For instance, consider a PoP with two possible operating configurations corresponding to High and Low capacity, with High/Low capacity states entered when traffic goes above/below a given threshold. Since the number of devices in sleep mode is larger in low capacity state configuration, the overall energy consumption decreases. This strategy can be replicated for more than two configurations. In this work,

we consider in details the cases for which PoPs can operate in two or three different operating configurations.

In the following, similarly to the energy to traffic proportional technologies, we assume normalized traffic values and normalized energy consumption. In particular, for simplicity, we assume that when the PoP is in an operating configuration in which the maximum manageable traffic is a fraction  $\rho$  of the full traffic capacity, the consumption of the PoP is also equal to fraction  $\rho$  of the consumption at full capacity. Thus, the energy savings can be computed using also in this case Eq. 5.7.

In the case of resource consolidation with *two operating configurations*, the PoP can support in one configuration the maximum traffic load and consumes the maximum amount, while, in the other, the PoP can process a limited amount of traffic and it has a lower energy consumption.

Then, when the traffic is above a fraction  $m$  of the peak traffic, the PoP is fully operative and consumes the maximum energy, that is normalized to 1. When traffic is below a fraction  $m$  of the peak, the PoP can work at a fraction  $m$  of its full capacity, and it is also consuming a fraction  $m$  of the consumption at full capacity. Thus,

$$e(\rho) = \begin{cases} m & \rho < m \\ 1 & \rho \geq m \end{cases} \quad (5.8)$$

Finally, we extend the previous case and assume that a PoP work in *three operating configurations*: one for low traffic, one for medium and one for maximum traffic. Thus,

$$e(\rho) = \begin{cases} l & \rho < l \\ m & l \leq \rho < m \\ 1 & \rho \geq m \end{cases} \quad (5.9)$$

where  $l$  and  $m$  are the traffic thresholds and corresponding energy consumption values for low and medium traffic, respectively.

The scenarios are sketched in Fig. 5.14, where the solid line represents the traffic and the dashed line the consumption.

## 5.4.2 Sleep mode policies at the ADSL modems

Since the traffic analysis, reported in Sec. 5.3, shows that ADSL users are typically inactive for long time periods, it may be advantageous to implement sleep mode policies in ADSL modems. Thus, when no or little traffic is present, modems can enter a *sleep mode* state with reduced energy requirements where no or limited transmission/reception capabilities are present. Indeed, current studies have shown that ADSL modems are consuming the same amount of energy independently from the fact that they are transmitting

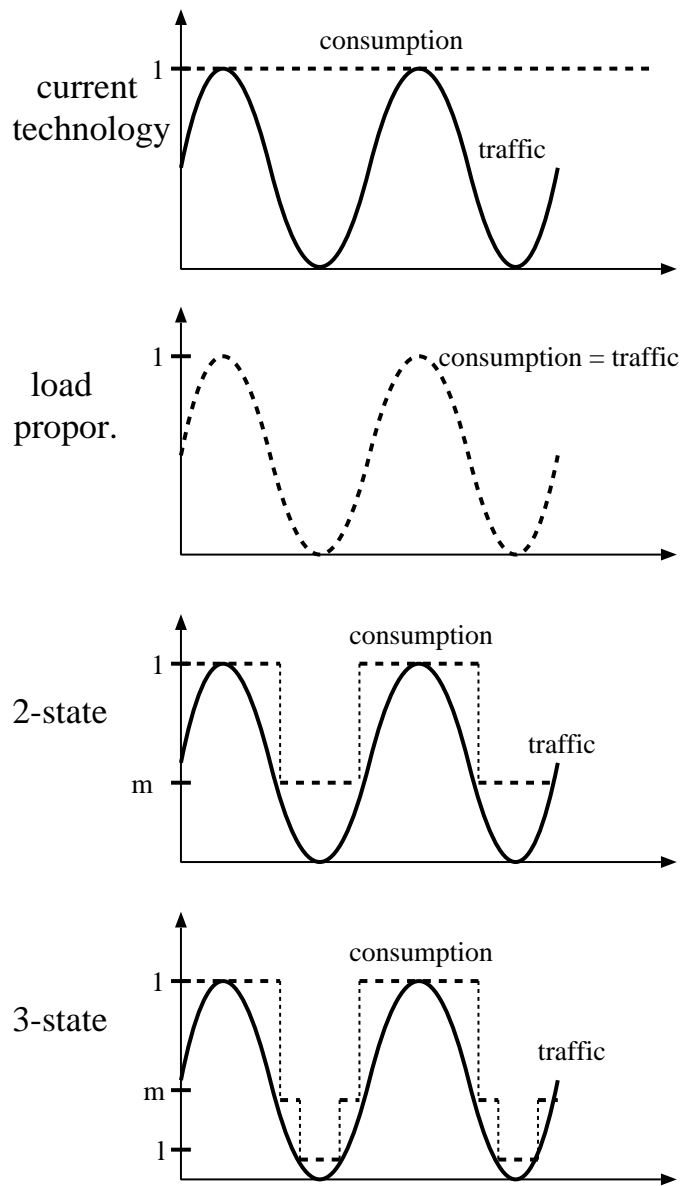


Figure 5.14: Energy savings schemes

or not [13,58]. Thus, all the time, that a modem is on without transmitting, represents an energy waste.

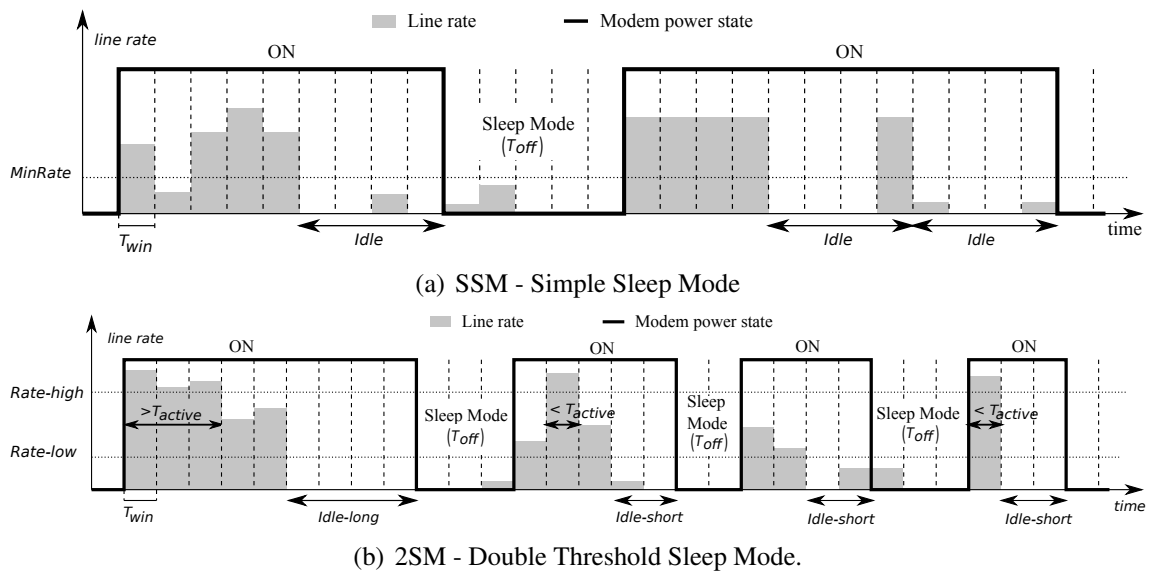


Figure 5.15: Example of evolution of the line rate and the modem power consumption obtained using two sleep mode policies.

ADSL2 standard [59] already incorporates recommendations about sleep mode policies, but at the moment no commercial modems are implementing them. Other studies have already proposed some solution to increase the energy efficiency in the ADSL access devices. In [60], authors propose a “roaming” scheme that allows end users traffic to be carried by nearby Wi-Fi channels when the ADSL/Wi-Fi link of the user enters sleep mode. In [61], energy savings obtained from the implementation of the ADSL2 standards are evaluated. In this case, no coordination among modems is required and each physical link is managed separately. Unfortunately, the study is limited to a single test case.

In this work, we consider two possible sleep mode state: in the first one, when entering sleep mode, both the DSLAM and users’ modem are completely turned OFF, i.e., the physical channel is lost; in the second one, the modems are able to enter a “low rate / low power ” state in which the physical link is working at a low bitrate in the order of few kb/s. with a reduced energy consumption. These solutions imply some technical details related to the modems functioning. For example, in the first case the physical channel is lost, and no traffic can be sent/received. This means that re-synchronization delays are experienced each time the line has to be woken up from the sleep mode.

As already stated before, characterizing the implementations details of the sleep mode policies is out the scope of this work. Those are strictly related to the modem hardware and technological constraints. We just assume that the modems have enough buffering resources to compensate the variations of the line rate introduced by the sleep mode policies.

In this work, since the line rate measurements are performed each time period of  $T_{win} = 10$  seconds, we evaluate the sleep mode policies according to this granularity in a postprocessing step and not in real time.

### Simple Sleep Mode (SSM) policy

we start considering a simple policy that we called Simple Sleep Model (SSM) which exploits a single line rate threshold, namely *MinRate*, to identify when the modem can enter the sleep mode.

In Fig. 5.15(a) is depicted the strategy reporting an example of line rate evolution over time. The modem monitors the line activity measuring the line rate. If the rate does not exceed *MinRate* for *Idle* seconds, then the line is assumed to be inactive and the modem enters sleep mode. As soon as the line rate is above *MinRate*, the modem has to be woken up and to become fully operative in a very short time.

### Double threshold Sleep Mode (2SM) policy

the SSM policy is very conservative, indeed, the modem can be woken up for traffic events not related to user's activity, e.g. background traffic related to control information. These events usually last for very short time and typically do not require high bitrate. Still, the modem has in any case to wait for *Idle* seconds before enter sleep mode again.

Therefore, it may be better that the modem, after these events, re-enters quickly the sleep mode state. The Double Threshold Sleep Mode (2SM) policy, thus, considers the evolution of the rate in  $T_{active}$  seconds and it exploits two traffic thresholds to better distinguish between users' activity and background traffic.

Fig. 5.15(b) shows an example of line rate evolution controlled by the 2SM policy. Similarly to before, as soon as the rate is higher than *Rate-low*, the modem is turned ON. Another threshold, *Rate-high*, is used to decide for how long the modem has to wait before entering the sleep mode. If the line rate is above *Rate-high* for at least  $T_{active}$  seconds, then the user is assumed to be active and the sleep mode is delayed waiting for an *Idle-long* inactivity period. This means that the modem is left ON for *Idle-long* seconds even if the rate goes below *Rate-high* in a second moment. Otherwise, we assume that the traffic is more likely to be related to control information than user's actions and a fast sleep mode idle time,  $Idle-short \leq Idle-long$ , is then used to enter sleep mode.

## 5.5 Results

The results of the previously described strategies are presented in this section. Remember that the evaluation of the achievable energy savings has been performed on real traffic and energy data. Thus, the energy savings, resulting from our evaluation, can be really obtained if the proposed strategies can be implemented in practice.




Figure 5.16: Energy savings for PoPs with linear proportional energy consumption




Figure 5.17: Energy savings for PoPs with cubic proportional energy consumption

### 5.5.1 Energy proportionality in the PoPs

In the following, the energy savings are computed considering the measured values of traffic which have been introduced in Sec. 5.1. For each PoP  $X$ ,  $e(\rho)$  in Eq. 5.5 is fed with  $\rho = \rho_X(n)$  as from Fig. 5.9.

#### Energy to traffic proportional technologies

the energy savings that can be achieved by using linear and cubic energy proportional devices are reported, respectively, in Fig. 5.16 and in Fig. 5.17 for various values of the parameter  $\alpha$  in (5.5). When  $\alpha = 0$ , the consumption is not proportional at all. Therefore,




Figure 5.18: Energy savings for PoPs with resource consolidation and two operating configurations




Figure 5.19: Energy savings for PoPs with resource consolidation and three operating configurations

in both cases no savings are possible and the consumption is the same as it is with current devices. As the quantity of energy that is proportional to load increases, savings are possible: when consumption is fully load proportional ( $\alpha = 1$ ), up to 50-60% of energy can be saved in the linear case, while in the cubic case energy savings can reach percentages between the 80 and 90%.

The savings depend on the traffic pattern, and, since traffic patterns are very similar in the PoPs, savings are almost the same for the different PoPs.






Figure 5.20: Percentage of energy savings for PoP  $B$  with resource consolidation and three operating configurations

### Resource consolidation

the energy savings achieved by each PoP, when resource consolidation with two operating configurations is used, are shown in Fig. 5.18. In the plot, the threshold value  $m$  varies from 0 to 1.

When  $m$  is small (below 0.2 of the peak traffic) no savings are possible, since traffic rarely drops below  $m$  and low capacity configuration is rarely entered. When  $m > 0.8$ , only the low capacity configuration is used, and savings depend on the actual value of  $m$  only. Depending on the traffic characteristics of the PoP, the optimal value of  $m$  slightly changes, but it is typically between 0.65 and 0.75. As shown in [62], the optimal value depends on the amount of energy that is saved in low capacity scheme and on the duration of the periods in which traffic is low and low capacity configuration can be entered.

The case of resource consolidation with three configurations is made in Fig. 5.19 for the case in which  $l = m/2$ . As expected, savings increase with respect to the two configurations case, but only marginally. Fig. 5.20 reports the savings for any value of  $l$  and  $m$ , with  $l \leq m \leq 1$ , for PoP  $B$  only. The optimal values of  $m$  and  $l$  are also in this case different for each PoP, but we evaluated, through an exhaustive search, that these values are close to  $m = 0.7$  and  $l = 1/2m$  for all the PoPs.

### 5.5.2 Sleep mode policies at the ADSL modems

The evaluation of the effectiveness of the sleep mode policies in the ADSL modems has been based on the following metrics:

- $T_{off}(i, l)$ , the duration of the  $i$ -th sleep mode interval for the line  $l$ ;




Figure 5.21: CDF all  $T_{off}$  periods for all the lines. Results refer to SSM with  $MinRate = 1$  kb/s

- $T_{off-total}(l)$ , the total time spent in the sleep mode by the line  $l$  corresponding to  $T_{off-total}(l) = \sum_i T_{off}(i,l)$ ;
- $N_{trans}(l)$ , the number of transitions from the sleep mode to the ON state occurred for a line  $l$ ;

Fig. 5.21 shows the CDF of all  $T_{off}$  time periods for all the lines considering SSM with  $MinRate = 1$  kb/s. Different patterns refer to different values of  $Idle$ . It is possible to notice that most of the periods lasts for just few hours. More in details, 80% of the periods in which the line is not active are below 6 hours. Note that shortening  $Idle$  increases the number of short  $T_{off}$  periods, since, for smaller  $Idle$ , it is possible to enter sleep mode state more frequently, but the average duration of the  $T_{off}$  period is smaller. This means that modem enters the sleep mode even when there is just a pause of the user activity. Thus, the modem does not wait for enough time before going into sleep mode state. Varying the  $MinRate$  we obtain similar CDFs.

The CDF of  $T_{off-total}$  for SSM is depicted in Fig. 5.22. Also in this case we show the results just for  $MinRate$  equal to 1 kb/s since for others rate the distributions are similar. We can appreciate that half of the lines can spend in sleep mode state more than 16 hours considering a conservative  $Idle$  equal to 30 minutes. Selecting more aggressive settings of  $Idle$  leads to higher savings. For instance, for  $Idle$  equal to 5 minutes, half of the users are staying in the sleep mode state more than 20 hours, since transitions to sleep mode are more frequent. These results show that the implementation of sleep mode strategies can be effective for most of the users.

In Fig. 5.23, we report the average time spent by users in the sleep mode state, that is

$$E[T_{off-total}(l)] = \frac{\sum_{i=1}^N T_{off-total}(l,i)}{N} \quad (5.10)$$




Figure 5.22: CDF of  $T_{off-total}$  for all the lines. Results refer to SSM with  $MinRate = 1$  kb/s.

where  $N$  is the total number of users of our data set.

To see the effect of background control traffic, we consider different values of  $MinRate$ . Recall that a modem can enter in sleep mode only if the line rate has not exceeded  $MinRate$  for  $Idle$  seconds. In case no  $MinRate$  is applied (e.g.,  $MinRate = 0$  kb/s) this corresponds to delay entering in sleep mode for any traffic events. Thus, the average time spent in the sleep mode state is very limited. This can be clearly seen in Fig. 5.23. As  $Idle$  is higher than 15 minutes,  $E[T_{off-total}]$  is almost zero, meaning that control traffic is generated periodically and idle time periods are not long enough to permit the modem to enter sleep mode.

Instead, for  $MinRate$  greater than zero, even a low value such as 1 kb/s brings significant improvements. Indeed, in the worst case with  $Idle$  equal to 30 minutes, we have  $E[T_{off-total}]$  equal to 15 hours. Decreasing  $Idle$  it is possible to gain up to 3 hours of sleep mode considering  $MinRate$  equal to 1 kb/s. If we further increase  $MinRate$ , no significant advantages are introduced; indeed, comparing 1 kb/s and 10 kb/s thresholds, the latter achieves a  $E[T_{off-total}]$  about 2 hours longer in the best case. Thus, the value of  $MinRate$  is not strongly affecting the time spent in the sleep mode state, provided it is larger than zero.

As stated in Sec. 5.3, each user can have a very different behavior concerning the traffic activity. This is confirmed considering the standard deviation of the silent periods  $std[T_{off-total}]$  reported in Fig. 5.24. The high values show that users are generating traffic in very different ways. However, results show that, on average, sleep mode state can last for hours. If some lines are staying in the sleep mode state for small periods of time or they do not enter in it at all, it means that sleep mode strategies are not effective just for these users.




Figure 5.23: Comparison of the average duration of total silent time  $T_{off-total}$  using the SSM policy with different settings.




Figure 5.24: Comparison of the standard deviation of the duration of total silent time  $T_{off-total}$  using the SSM policy with different settings.

### Impact on latency

the number of transitions from sleep mode to ON state is an important metric to evaluate the QoS of the lines. In fact, each transition imply a delay needed to re-synchronize the line signals or to allow the line to reach the maximum rate. Thus, it is better to choose wisely the settings in order to limit the number of these transitions. In Fig. 5.25 we report the average number of state transitions  $E[N_{trans}]$  for different values of  $MinRate$  and  $Idle$  considering the SSM policy. In case no  $MinRate$  is applied, the number of transitions is very large for low values of  $Idle$ . Indeed, the line continuously switches from one state to the other. Instead, the number of state transitions decreases considerably for  $T_{idle} \geq 20$  minutes since the modems do not enter sleep mode anymore. Considering values of




Figure 5.25: Comparing the number of transitions  $N_{trans}$  from sleep to ON state using the SSM policy with different settings.

*MinRate* equal to 1 kb/s or higher, we can see that there are not significant differences. In particular, we can see that the number of transitions is decreasing until *Idle* equal to 15 minutes, while for larger values of *Idle* the number of transitions is almost stable around 5. Thus, choosing a *Idle* of 15 minutes can be the best trade-off between the number of transitions and the time spent in the sleep mode state for the SSM policy.

Furthermore, in Fig. 5.26, which shows the CDF of  $N_{trans}$  for *MinRate* equal to 1 kb/s, it is possible to notice that most of the users experiences a limited number of transitions. Indeed, considering  $T_{idle}$  equal to 15 minutes, the 80% of users has 10 or less transitions during the day. Thus, only few users are switching from the sleep mode to the ON state a very large number of time. As expected, decreasing  $T_{idle}$ , the number of transitions experienced by the 80% of users increases to 25 transitions, while with  $T_{idle}$  set to 30 minutes, the same percentage of users has at most 7 transitions. Similar results are retrieved for higher values of *MinRate* without any significant change.

## 2SM improvements

considering the 2SM policy, the evaluation of  $E[T_{off-total}]$  is shown in Fig. 5.27 for different settings. Given the previous considerations, we consider *Rate-low* = 1 kb/s, *Idle-long* = 15 min and we set  $T_{active}$  equal to 2 min.

We notice that the values of  $E[T_{off-total}]$  are almost constant with respect to *Rate-high* but decreasing *Idle-short* it is possible to spend more time in the sleep mode state. In particular, using *Rate-high* = 100 kb/s and *Idle-short* = 2 min, a line is silent for 18.2 hours in average. Instead, using *Rate-high* = *Rate-low* = 1 kb/s and *Idle-short* = *Idle-long* = 15 min, i.e., the SSM policy,  $E[T_{off-total}]$  accounts for 16.3 hours. This means that with more aggressive settings it is possible to save more energy. However, these savings have to be compared with the number of state transitions needed as reported in Fig. 5.28.




Figure 5.26: CDF of the number of transitions  $N_{trans}$  from sleep to ON state using the SSM policy with  $MinRate = 1kb/s$  and different settings of  $Idle$ .




Figure 5.27: Comparison of the average duration of total silent time  $T_{off-total}$  using the 2SM policy with  $Rate-low = 1 kb/s$  and  $T_{active} = 2 min$ .

In fact, we can notice that this aggressive setting introduces 8 more state transitions with respect to the SSM case. This means that the 2 hours of benefits obtained can imply some performance impairments and more conservative settings could represent a better performance compromise.

### 5.5.3 Achievable energy reduction and monetary savings

In the following, we compare the energy savings that can be achieved by the different strategies and we quantify approximately which can be the energy saved in one-year long period for each PoP and how large can be the correspondent monetary savings.




Figure 5.28: Comparing the number of state transitions  $N_{trans}$  using the 2SM policy with  $Rate-low = 1$  kb/s and  $T_{active} = 2$  min.

### Energy savings at the PoPs

the optimal performance of each proposed strategy for the PoPs is shown in Fig. 5.29. For the linear and cubic energy proportional technologies, we report the results for  $\alpha = 1$  since this value guarantees the maximum possible savings. Moreover, it may be considered as an upper bound of the energy savings that can be obtained with energy proportional schemes.

For the resource consolidation schemes, instead, for each PoP we perform an exhaustive search in order to find which are the best values of the thresholds in the case of the two and of the three operating configurations. In the figure, the PoPs are ordered by increasing number of networking devices; notice that the size of the PoPs does not influence the possible energy savings.

As expected, cubic proportional technology achieves the highest savings with more than 80% of energy saved for any PoP. Still consistent are the savings obtained with linear proportionality, which are typically about or higher 50%. Even if a lot of effort is spent to develop network devices with energy proportionality characteristics, it will require some time before that they will become available on commercial devices. Instead, resource consolidation strategies are likely to be deployable in the near future and, even if lower, the possible savings can be quite large, from 20-30% of the two configurations case to around 40% with three configurations levels. These solutions are probably very promising as short-term relief to high ISPs electricity bills. Moreover, note that between two and three configurations the energy savings present a not significant difference, thus it may be sufficient to just implement two operating configurations to achieve interesting savings.




Figure 5.29: Energy savings for the three models




Figure 5.30: Average energy consumption per users. SSM with  $MinRate = 1$  kb/s and  $T_{idle} = 15$  min. 2SM with  $Rate-low = 1$  kb/s,  $Rate-high = 100$  kb/s,  $Idle-short = 5$  min,  $Idle-long = 15$  min and  $T_{active} = 2$  min.

### Energy savings at the ADSL modems

we evaluate the energy consumption for the two proposed sleep mode policies and we compare them with respect to the case in which the modems are always ON. We assume that the power consumption of a DSLAM modem port is around 1 W while an users' modem consumes about 5 W according to the values reported in [58, 60].

As introduced in Sec. 5.4, the power consumption of modems in the sleep mode state depends strongly on their implementation. In this study we consider two possible cases of the sleep mode state: (i) the modem is switched OFF, thus consuming approximately 0 W,



(ii) the modem can vary the line rate reducing it to save energy. Based on the considerations reported previously, we select the following settings for the policies which represent the best trade-off between time spent in the sleep mode state and number of transitions: for SSM  $MinRate = 1$  kb/s and  $Idle = 15$  min, for 2SM  $Rate-low = 1$  kb/s and  $Rate-high = 100$  kb/s, while  $Idle-short = 5$  min,  $Idle-long = 15$  min and  $T_{active} = 2$  min.

Using these settings we obtained the energy consumption reported in Fig. 5.30. The x-axis reports the ratio between the energy consumed by the modem in the sleep mode state and the energy consumed in the ON state. Thus, when the x-coordinate is equal to 0, it is the case in which the modem is powered OFF when it enters sleep mode. While for larger values, it indicates the fraction of energy that the modem consumes in the sleep mode with respect to the ON state. For instance, when this ratio is equal to 0.5 it means that in that scenario the consumption of the modem during the sleep mode is halved with respect to the ON state.

In the figure, there are reported for comparison the amount of energy consumed by modems for the ON state (i.e., the modems are always on), for the SSM and for the 2SM policies. Note that the energy is evaluated with two different scales. Left y-axis reports the average amount of energy consumed per user in one day considering the average values obtained from analysis of the dataset. Instead, right y-axis reports the energy consumed in one year assuming 1 million of customers. These results are obtained assuming the monitored users are representative for a wider population and power consumption is similar in different days. These are quite strong assumptions but they allow to easily have a qualitative indication of the possible gain of sleep mode policies when widely adopted.

We can notice that the energy consumed by the two strategies is very similar, with 2SM performing a little better. Indeed, the decrease of the energy consumption guaranteed by the simpler SSM policy is already astonishing. Considering an intermediate scenario, in which we set the power consumption of modems in sleep mode equal to 0.5, we have that the SSM allows to easily save 35%, which consists to save 18 GWh in one year for a population of 1 million users. Note that, since the DSLAM modems accounts for 1/6 of the total energy consumption, an operator that is willing to implement a sleep mode strategy only at DSLAM can save at least 3 GWh, which corresponds to about 330 kEuro per year, assuming a price of electricity equal to 0.0936 Euro/kWh [32].

If we consider the best scenario (i.e., the modem does not consume during the sleep mode), the energy saved is about 38 GWh, which is larger, but not in a significant manner. This means that it is possible to achieve huge savings also considering to not switch off completely the modems, but just using a “low power” state.

### **Energy and monetary savings at the PoPs in one-year long period**

we quantified which can be the real energy and monetary savings implementing energy efficient policies at the PoPs that we take into account, considering also to apply sleep mode policies at the DSLAMs associated to these PoPs.

In order to get as much as possible realistic savings, we choose the more conservative policies. For what concerns the PoPs, we estimate the savings in the case of the resource consolidation strategy with two operating configurations. The value of threshold  $m$  has been set equal to the optimal one, as in Fig. 5.29. We assume that the energy consumed in the High operating configuration corresponds to the current consumption of networking devices, thus  $e(1) = \min(\mathcal{E})$  as from Table 5.2. For the ADSL modems, we select the SSM policy and we use the same setting of the parameters employed to retrieve the results of Fig. 5.30.

In Table 5.3, there are reported in the columns, respectively, the energy savings at each PoP, the energy savings in the DSLAMs associated to a PoP, the total energy savings that can be achieved in a PoP and the correspondent total monetary savings. The price of electricity is set again equal to 0.112 Euro/kWh.

Note that, even if the energy savings percentages are almost equal for all PoPs, just in the larger PoPs the amount of saved energy, and of money, is substantial. This means that only in these PoPs there is an economical convenience to invest in the implementation of energy efficient policies.

Looking at the energy saved by applying the SSM policy at the DSLAMs modems, it possible to notice, in the PoPs with a large number of ADSL users, that the saved energy is significant even if the DSLAMs modems consume only about 1 W.

Table 5.3: Energy savings at the PoPs

ID	$S_{PoP}$ [MWh]	$S_{ADSL}$ [MWh]	$S_{ToT}$ [MWh]	Money savings [kEuro]
<i>A</i>	8.5	0	8.5	0.95
<i>B</i>	153.8	103.8	257.7	28.96
<i>C</i>	85.9	67.8	153.8	17.2
<i>D</i>	60	7	67.1	7.5
<i>E</i>	125	3.3	128.3	14.4
<i>F</i>	352.3	165.6	517.87	58
<i>G</i>	195.4	174.2	369.64	41.4
<i>H</i>	103.7	67.1	170.84	19.1

## 5.6 Conclusions

We exploited a wide dataset of real measurements, collected in the operative network of FASTWEB, a national-wide Italian operator, to perform an extensive analysis of the energy and traffic of the PoPs and to characterize the traffic profiles of the ADSL lines. We observed that PoPs present a traffic profile that is varying in time, while the consumed energy is not dependent on the traffic load. In particular, we found that in the PoPs there is a correlation between the consumed energy and external temperature, mainly due to

the air conditioning system. This correlation is especially strong in the PoPs that employ air conditioners with the free cooling option. These results motivated us to estimate the energy savings that can be possibly achieved with policies that can adapt the energy consumption to the traffic load currently managed by the PoPs.

We considered as possible strategies at the PoPs to implement solutions based on energy-to-traffic proportional technologies or on the resource consolidation practice. We estimated that the implementation of these policies brings significant energy savings; using the more conservative strategy it is possible to save about 30% of energy.

The traffic analysis of the ADSL lines reports that the lines are usually active for limited periods of time or present a very low bit rate. Thus, we considered to implement two sleep mode policies that, according to the traffic activity, can decide to put the modems at DSLAM and at user side in a sleep mode state characterized by a reduced power consumption and transmitting capabilities. We evaluated that, even if each modem consumes just few Watts, the overall achievable energy savings in one year can be in the order of tens of GWh since the total number of these devices in the access networks is very high.

The obtained results thus indicate the effectiveness of the considered energy efficient policies to reduce energy and the monetary costs related to it. Our conclusion is that the implementation of these energy efficient policies is of strategical importance for the ISPs.



# Chapter 6

## Conclusions and Future Work

In this thesis, we tackled the issue of energy efficiency in wired telecommunication networks proposing energy-efficient strategies for each network segment. We evaluated the achievable energy savings and we estimated which is the impact of these strategies also over other factors, such as the QoS and the CapEx.

We started by proposing a power-aware design for the LTs of WR networks. We introduced a MILP to solve the design problem in an optimal way. We compared the solutions retrieved for two different target optimizations: power and CapEx minimization. We found that the power-aware solutions can effectively lead to network configurations with a reduced consumed energy, but at the cost of an increase of the CapEx. Anyway, we evaluated that the economical savings, achieved thanks to the the reduced energy consumption, can recover the additional investments. Furthermore, within this activity, we develop some heuristics that well approximate the optimal results.

Next, we estimated the advantages of a dynamic configuration of WR networks. Indeed, since traffic is varying according a daily pattern, we quantified, over several different traffic scenarios and networks the energy savings achievable from a multi-period configuration of the networks. In details, we developed three heuristics to solve this problem targeting the joint minimization of power consumption and reconfiguration costs, that are represented by the amount of reconfigured traffic. Indeed, between two consecutive time periods, the network has to be reconfigured and a large amount of traffic reconfigured may penalize the QoS perceived by the users. Results show that the heuristics can effectively reduce the energy consumption of the networks and, in the meanwhile, they limit the reconfiguration costs. Significant energy and economical savings can be thus achieved considering the multi-period configuration with respect to a static network design.

Further solution for the core networks, it is the introduction of more energy efficient switching paradigms. Within this activity, we thus evaluated the energy efficiency time-domain optical sub-wavelengths switching technologies with respect to legacy circuit switching solutions, which are widely employed nowadays. Through simulations,

we retrieved the energy consumption in a simple network case for three different sub-wavelength switching technologies and for circuit switching paradigms. We retrieved that, in low traffic scenarios, sub-wavelength switching technologies can effectively reduce the consumption on average of about 30%.

For what concerns aggregation and access networks, we focused respectively on the PoPs and on ADSL modems. In particular for the PoPs, (i) we performed a characterization of the energy and traffic profiles of the PoPs and (ii) we proposed and evaluated possible energy saving strategies, according to the findings retrieved with the characterization. The profiling of the energy consumed by the PoPs showed that the consumption is correlated to the external temperature, while it is totally independent from the traffic processed. This result motivated the choice to investigate some strategies that can make the energy consumption proportional to the traffic processed by the PoP. Two different approaches have been thus considered. The first one assumes that there is a direct proportionality between the energy consumed and the traffic processed. The second one is based on resource consolidation, in which a PoP can use different operating configurations. In each of them, the PoP can manage a specific traffic load with a given energy consumption. The achievable savings have been evaluated using the traffic measures provided by FASTWEB, a main Italian operator, and results show that savings up to 30% are likely to be achievable.

As possible strategy for the ADSL modems, since the analysis of the activity of ADSL users shows that the majority of them is not generating traffic for most of the time, we proposed to reduce the energy consumption in access networks implementing sleep mode policies at the ADSL modems. Indeed, the modem consumes the same if it is transmitting or not. Thus, when user is not active, the modem may be enforced to a sleep mode state in which its transmission capabilities are limited and the power requirements reduced. In this activity, two different sleep mode approaches have been investigated: the first one presents a conservative behavior, while the other one is more aggressive and it tries to put the modem in sleep mode more frequently. Results show that significant energy savings are possible in both cases. Assuming a conservative setting of the parameters, it is possible to reduce the energy consumption of ADSL modems of about 35% with respect to the case of not having implemented sleep mode policies. Furthermore, the aggressive approach does not bring any advantages. Indeed, there are no particular improvements, for what concerns the energy saved, and in the meanwhile the aggressive approach penalizes the QoS perceived by the users, due to more frequent transitions from the sleep mode to the on state.

In these three years of research, we investigated how to improve the energy efficiency of telecommunication networks. Different strategies have been analyzed and we saw that consistent energy savings can be achieved in every network segment. However, before being able to implement these strategies in real networks, it is required to solve some technical issues.

For example, in the multi-period configuration of core networks it is necessary to

---

define how to manage the reconfiguration of the networks. Communication protocols, which exchange information about the reconfiguration of the network, should be developed. Routers should sustain the possibility to switch off some of their modules and each component should be designed in order to not suffer from the switching off and on. Similar examples can be the implementation of a sleep mode policy in an ADSL modem or the application of the resource consolidation practice to a PoP.

Some of these issues are likely to be solved easily, but it is necessary to encourage manufacturers, network operators, ISPs and all the actors that are all involved in the design, management and use of the telecommunication networks, to invest in this direction.





# Bibliography

- [1] C. Lange, D. Kosiankowski, C. Gerlach, F. Westphal, and A. Gladisch, “Energy Consumption of Telecommunication Networks,” in *Proc. of the ECOC, Vienna, Austria*, September 2009.
- [2] E. Le Rouzic, “Network evolution and the impact in core networks,” in *Optical Communication (ECOC), 2010 36th European Conference and Exhibition on*, sept. 2010, pp. 1–8.
- [3] M. Webb, “SMART 2020: Enabling the Low Carbon Economy in the Information Age,” The Climate Group. London, June 2008.
- [4] Global Action Plan, “An Inefficient Truth,” Global Action Plan Report, <http://globalactionplan.org.uk>, December 2007.
- [5] S. Pileri, “Energy and communication: engine of the human progress,” INTELEC 2007 keynote, Rome, September 2007.
- [6] M. Etoh, T. Ohya, and Y. Nakayama, “Energy consumption issues on mobile network systems,” in *Proceedings of 2008 International Symposium on Applications and the Internet (IEEE SAINT 2008)*, July 2008.
- [7] International Energy Agency Web Page, “<http://www.iea.org/journalists/fastfacts.asp>,” 2010.
- [8] Towards Real Energy-efficient Network Design (TREND), <http://www.fp7-trend.eu>, 2010.
- [9] ECONET Project, funded by the European 7th Framework Programme, “<http://www.econet-project.eu>.”
- [10] GreenTouch Consortium, “<http://www.greentouch.org/>.”
- [11] D-Link Green Products, “<http://www.dlinkgreen.com/greenproducts.asp>.”
- [12] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, “Energy Consumption in Optical IP Networks,” *J. Lightwave Technol.*, vol. 27, no. 13, pp. 2391–2403, July 2009.
- [13] A. Valenti, F. Matera, and G. Beleffi, “Power consumption measurements of access networks in a wide geographical area test bed and economic perspectives,” in *Future Network Mobile Summit (FutureNetw)*, 2012, july 2012.
- [14] 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.
- [15] I. Cerutti, L. Valcarengi, and P. Castoldi, “Designing Power-Efficient WDM Ring Networks,” in *Networks for Grid Applications*, ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, A. Doulamis, J. Mambretti, I. Tomkos, and T. Varvarigou, Eds. Springer Berlin Heidelberg, 2010, vol. 25, pp. 101–108.
- [16] J. Restrepo, C. Gruber, and C. Machuca, “Energy profile aware routing,” in *Communications*

- Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*, june 2009, pp. 1–5.
- [17] E. Bonetto, L. Chiaraviglio, D. Cuda, G. Gavilanes Castillo, and F. Neri, “Optical technologies can improve the energy efficiency of networks,” in *Optical Communication, 2009. ECOC '09. 35th European Conference on*, sept. 2009, pp. 1–4.
- [18] G.Griffa, “Ngn2 and energy opex savings,” Torino, Italy, November 2008.
- [19] Colfax International, “<http://www.colfaxdirect.com/>.”
- [20] Fiber Optics Inventory Store, “<http://www.lightwavestore.com/>.”
- [21] A. Nag and M. Tornatore, “Optical network design with mixed line rates,” *Optical Switching and Networking*, vol. 6, no. 4, pp. 227 – 234, 2009, special Issue: Selected Papers from the Second International Symposium on Advanced Networks and Telecommunication Systems (ANTS 2008).
- [22] 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.
- [23] A. Ahmad, A. Bianco, E. Bonetto, D. Cuda, G. G. Castillo, and F. Neri, “Power-aware logical topology design heuristics in wavelength-routing networks,” in *Proc. of the ONDM, Bologna, Italy*, February 2011.
- [24] 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.
- [25] F. Idzikowski, L. Chiaraviglio, and E. Bonetto, “Ewa: An adaptive algorithm for energy saving in ip-over-wdm networks,” in *Networks and Optical Communications (NOC), 2012 17th European Conference on*, june 2012, pp. 1–6.
- [26] L. Chiaraviglio, M. Mellia, and F. Neri, “Minimizing ISP Network Energy Cost: Formulation and Solutions,” *IEEE/ACM Transactions on Networking*, July 2011.
- [27] 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*, vol. 8, no. 3, pp. 181–200, July 2011.
- [28] “Sndlib: library of test instances for survivable fixed telecommunication network design,” <http://sndlib.zib.de/home.action>, February 2012.
- [29] “Cisco CRS-1 Production Brochure,” [http://www.cisco.com/en/US/prod/collateral/routers/ps5763/prod\\_brochure0900aecd800f8118.pdf](http://www.cisco.com/en/US/prod/collateral/routers/ps5763/prod_brochure0900aecd800f8118.pdf), October 2008.
- [30] 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.
- [31] 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, 2008.
- [32] “Eurostat database web page,” <http://epp.eurostat.ec.europa.eu/>, February 2012.
- [33] M. Yamada, T. Yazaki, N. Matsuyama, and T. Hayashi, “Power efficient approach and performance control for routers,” in *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*, june 2009, pp. 1–5.

- [34] 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.
- [35] M. Zhang, C. Yi, B. Liu, and B. Zhang, "GreenTE: Power-aware traffic engineering," in *Proc. of the 18th IEEE International Conference on Network Protocols (ICNP)*, October 2010, pp. 21–30.
- [36] A. Farrel and I. Bryskin, *GMPLS Architecture and Applications*. Morgan Kaufmann Publishers, 2006.
- [37] C. Qiao, "Labeled optical burst switching for ip-over-wdm integration," *Communications Magazine, IEEE*, vol. 38, no. 9, pp. 104–114, sep 2000.
- [38] P. Pavon-Marino and F. Neri, "On the myths of optical burst switching," *Communications, IEEE Transactions on*, vol. 59, no. 9, pp. 2574–2584, september 2011.
- [39] T. Legrand, B. Cousin, and N. Brochier, "Performance evaluation of the labelled obs architecture," in *Wireless And Optical Communications Networks (WOCN), 2010 Seventh International Conference On*, sept. 2010, pp. 1–6.
- [40] I. Widjaja and I. Saniee, "Simplified layering and flexible bandwidth with twin," in *Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture*, ser. FDNA '04. New York, NY, USA: ACM, 2004, pp. 13–20.
- [41] D. Chiaroni, G. Buforn, C. Simonneau, S. Etienne, and J.-C. Antona, "Optical packet add/drop systems," in *Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference, 2010 Conference on (OFC/NFOEC)*, march 2010, pp. 1–3.
- [42] T. Bonald, S. Oueslati, J. Roberts, and C. Roger, "Swing: Traffic capacity of a simple wdm ring network," in *Teletraffic Congress, 2009. ITC 21 2009. 21st International*, sept. 2009, pp. 1–8.
- [43] A. Ahmad, A. Bianco, E. Bonetto, D. Cuda, G. Castillo, and F. Neri, "Power-aware logical topology design heuristics in wavelength-routing networks," in *Optical Network Design and Modeling (ONDM), 2011 15th International Conference on*, feb. 2011, pp. 1–6.
- [44] Public deliverable D1.3 of TREND project, "<http://www.fp7-trend.eu/system/files/content-public/307-d13-final-report-ira-11/trend-d13.zip>."
- [45] JDSU, Tunable Multiprotocol XFP Optical Transceiver, "<http://www.jdsu.com/ProductLiterature/jxp01tmac1cx5-ds-oc-ae.pdf>."
- [46] B. Skubic, E. In de Betou, T. Ayhan, and S. Dahlfors, "Energy-efficient next-generation optical access networks," *Communications Magazine, IEEE*, vol. 50, no. 1, pp. 122–127, january 2012.
- [47] Oclaro, TL5000 ITLA, "[http://www.oclaro.com/datasheets/TL5000DCJ\\_Datasheet\\_-\\_D00255-PB\\_%5B02%5D.pdf](http://www.oclaro.com/datasheets/TL5000DCJ_Datasheet_-_D00255-PB_%5B02%5D.pdf)."
- [48] Ignis, Integrated tunable laser Assembly S7500, "[http://ignisblogg.com/wp-content/uploads/2011/01/S7500\\_Tunable\\_Laser\\_Module.pdf](http://ignisblogg.com/wp-content/uploads/2011/01/S7500_Tunable_Laser_Module.pdf)."
- [49] Emcore, ITLA (Integrable Tunable Laser Assembly), "[http://www.emcore.com/fiber\\_optics/telecom/itla](http://www.emcore.com/fiber_optics/telecom/itla)."
- [50] <http://tstat.polito.it>.
- [51] A. Finamore, M. Mellia, M. Meo, M. M. Munafò, and D. Rossi, "Experiences of Internet Traffic Monitoring with Tstat," *IEEE Network*, vol. 25, no. 3, pp. 8–14, 2011.

- [52] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsang, and S. Wright, "Power awareness in network design and routing," in *IEEE INFOCOM 2008*, Phoenix, USA, April 2008.
- [53] L. Liu, H. Wang, X. Liu, X. Jin, W. B. He, Q. B. Wang, and Y. Chen, "Greencloud: a new architecture for green data center," in *Proceedings of the 6th international conference industry session on Autonomic computing and communications industry session*, ser. ICAC-INDST '09. New York, NY, USA: ACM, 2009, pp. 29–38.
- [54] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, "Elastictree: saving energy in data center networks," in *Proceedings of the 7th USENIX conference on Networked systems design and implementation*, ser. NSDI'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 17–17.
- [55] J. Li, Z. Li, K. Ren, and X. Liu, "Towards optimal electric demand management for internet data centers," *Smart Grid, IEEE Transactions on*, March 2012.
- [56] D. Abts, M. R. Marty, P. M. Wells, P. Klausler, and H. Liu, "Energy proportional datacenter networks," *SIGARCH Comput. Archit. News*, vol. 38, pp. 338–347, June 2010. [Online]. Available: <http://doi.acm.org/10.1145/1816038.1816004>
- [57] R. Bolla, R. Bruschi, and M. Listanti, "Enabling backbone networks to sleep," *Network, IEEE*, vol. 25, no. 2, pp. 26–31, 2011.
- [58] A. O. Calchand, P. Branch, and J. But, "Analysis of Power Consumption in Consumer ADSL Modems," Centre for Advanced Internet Architectures, Swinburne University of Technology, Melbourne, Australia, Tech. Rep. 100125A, January 2010.
- [59] "G.992.3: Asymmetric Digital Subscriber Line transceivers 2 (ADSL2)." [Online]. Available: <http://www.itu.int/rec/T-REC-G.992.3/>
- [60] E. Goma and al., "Insomnia in the Access: or How to Curb Access Network Related Energy Consumption," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 338–349, August 2011.
- [61] A. Gladisch, C. Lange, and R. Leppla, "Power Efficiency of Optical versus Electronic Access Networks," in *Optical Communication, 2008. ECOC 2008. 34th European Conference on*, September 2008, pp. 1–4.
- [62] M. A. Marsan, L. Chiaraviglio, D. Ciullo, and M. Meo, "Optimal Energy Savings in Cellular Access Networks," in *Proceedings of the 1st International Workshop on Green Communications (GreenComm '09)*, Dresden, Germany, June 2009.