low Energy COnsumption NETworks

DELIVERABLE D2.1

END-USER REQUIREMENTS, TECHNOLOGY SPECIFICATIONS AND BENCHMARKING METHODOLOGIES

Grant Agreement Number:	**258454**
Project Acronym:	**ECONET**
Project Title:	**low Energy COnsumption NETworks**
Funding Scheme:	Collaborative Project
Starting Date of the Project:	01/10/2010 *dd/mm/yyyy*
Duration:	36 months
Project Coordinator:	Name: Raffaele Bolla Phone: +39 010 353 2075 Fax: +39 010 353 2154 e-mail: raffaele.bolla@unige.it

Due Date of Delivery:	M9 *Mx* (31/06/2011 *dd/mm/yyyy*)
Actual Date of Delivery:	28/09/2011 *dd/mm/yyyy*
Workpackage:	**WP2 – *User Requirements and Manufacturing Specifications***
Nature of the Deliverable:	R
Dissemination level:	PU
Editors:	CNIT, MLX, ALU, LQDE, TEI, TELIT, GRNET, NASK, ETY
Version:	1.0

List of the Authors

CNIT	CONSORZIO NAZIONALE INTERUNIVERSITARIO PER LE TELECOMUNICAZIONI
Raffaele Bolla, Roberto Bruschi, Alessandro Carrega, Franco Davoli, Paolo Lago, Chiara Lombardo	
ALU	ALCATEL-LUCENT ITALIA S.p.A.
Pasquale Donadio, Giorgio Parladori	
MLX	MELLANOX TECHNOLOGIES LTD - MLNX
Gil Bloch, Dvir Libhaber, Nir Strugo	
LQDE	LANTIQ Deutschland GmbH
Lorenzo Di Gregorio, Andrea Morici	
TEI	ERICSSON TELECOMUNICAZIONI
Massimo Enrico	
TELIT	TELECOM ITALIA S.p.A.
Flavio Cucchietti, Luca Giacomello, Gianluca Griffa, Paolo Salamandra, Diego Suino	
GRNET	GREEK RESEARCH AND TECHNOLOGY NETWORK S.A.
Vasiliki Giannikopoulou, Chrysostomos Tziouvaras, Constantinos Vassilakis, Anastasios Zafeiropoulos	
NASK	NAUKOWA I AKADEMICKA SIEC KOMPUTEROWA
Piotr Arabas, Przemyslaw Jaskola, Mariusz Kamola	
ETY	ETHERNITY NETWORKS LTD
David Levi	

Disclaimer

The information, documentation and figures available in this deliverable are written by the ECONET Consortium partners under EC co-financing (project FP7-ICT-258454) and do not necessarily reflect the view of the European Commission.

The information in this document is provided "as is", and no guarantee or warranty is given that the information is fit for any particular purpose. The reader uses the information at his/her sole risk and liability.

Copyright

The ECONET Consortium consists of:

CONSORZIO NAZIONALE INTERUNIVERSITARIO PER LE TELECOMUNICAZIONI,

ALCATEL-LUCENT ITALIA S.p.A.,

MELLANOX TECHNOLOGIES LTD - MLNX,

LANTIQ Deutschland GmbH,

ERICSSON TELECOMUNICAZIONI,

TELECOM ITALIA S.p.A.,

GREEK RESEARCH AND TECHNOLOGY NETWORK S.A.,

NAUKOWA I AKADEMICKA SIEC KOMPUTEROWA,

DUBLIN CITY UNIVERSITY,

TEKNOLOGIAN TUTKIMUSKESKUS VTT,

POLITECHNIKA WARSZAWSKA,

NETVISOR INFORMATIKAI ES KOMMUNIKACIOS ZARTKORUEN MUKODO RESZVENYTARSASAG,

ETHERNITY NETWORKS LTD,

LIGHTCOMM S.R.L.,

INFOCOM S.R.L.

Table of Contents

DISCLAIMER **3**

COPYRIGHT **3**

TABLE OF CONTENTS **4**

1 EXECUTIVE SUMMARY **6**

2 INTRODUCTION **6**

2.1 SCOPE OF THE DELIVERABLE 6
2.2 OBJECTIVES 7
2.3 STRUCTURE OF THE DOCUMENT 8

3 ENERGY CONSUMPTION IN NEXT-GENERATION NETWORKS: TRENDS AND TECHNOLOGIES **8**

3.1 INTERNET TOPOLOGY EVOLUTION 9
3.2 TRAFFIC VOLUMES AND PROFILES 10
3.2.1 Traffic statistics based on protocol breakdown 10
3.2.2 Access speeds 12
3.2.3 Traffic volumes and profiles 13
3.2.4 Traffic projections for the network of tomorrow 14
3.3 ENERGY CONSUMPTION AND ENERGY EFFICIENCY DEMAND 15
3.4 NEXT GENERATION NETWORKS TECHNOLOGIES 22
3.4.1 Optical backbone networks 22
3.4.2 Next generation technologies in the access part 24
3.4.3 Content centric networks 27
3.4.4 Autonomic networking and Cloud computing 28
3.5 ENERGY-AWARE DESIGN SPACE 29
3.5.1 Energy-aware design of the network device data plane functions 29
3.5.2 Energy-aware strategies design at the control plane 30
3.6 STANDARDIZATION BODIES 31
3.7 SUMMARY 33

4 ENERGY EFFICIENCY REQUIREMENTS FOR NEXT GENERATION NETWORKS AND DEVICES **34**

4.1 ENERGY CONSUMPTION IN TODAY'S NETWORK DEVICES 34
4.1.1 Access networks 35
4.1.2 Network terminations 44
4.1.3 Switching systems 48
4.1.4 IP routers 52
4.1.5 Network cards 55
4.1.6 Direct measurements in off-the-shelf devices 57
4.2 REFERENCE NETWORK SCENARIOS 61
4.2.1 Reference Telecom Operator scenario by TELIT 62
Reference Internet Service Provider scenario by GRNET 71
4.2.2 Reference Internet Service Provider scenario by NASK 86

4.3 ENERGY EFFICIENCY REQUIREMENTS AND TARGET ENERGY SAVING FOR NEXT GENERATION NETWORKS DEVICES 91
4.3.1 Energy saving requirements and targets for network terminations........ 95
4.3.2 Energy saving requirements and targets for access networks........ 100
4.3.3 Energy saving requirements and targets for metro/transport devices (L2/L2.5 switches) 103
4.3.4 Energy saving requirements and targets for core devices (IP router) 105
4.4 SUMMARY 107

5 TECHNOLOGY SPECIFICATIONS OF ENERGY EFFICIENT DEVICES 108

5.1 FUTURE ENERGY-AWARE TECHNOLOGIES ON THE NETWORK DEVICE DATA PLANE 109
5.1.1 Re-engineering 109
5.1.2 Dynamic adaptation 110
5.1.3 Sleeping/Standby 112
5.2 FUTURE ENERGY-AWARE MECHANISMS ON THE CONTROL PLANE 115
5.2.1 Energy aware traffic engineering 115
5.2.2 Resource allocation & Algorithmic efficiency 120
5.2.3 Virtualization techniques 121
5.3 ENERGY CONSUMPTION PROFILES FOR EACH NETWORK DEVICE ELEMENT 125
5.3.1 Energy consumption profiles of switch systems 125
5.3.2 Energy consumption profiles of VDSL home gateways 128
5.3.3 Energy consumption profiles of FPGA based systems 132
5.3.4 Energy consumption profiles of software routers 134
5.4 MONITORING AND MANAGEMENT ISSUES 137
5.5 SUMMARY 139

6 TARGET ENERGY SAVING AND GENERAL PERFORMANCE INDEXES 140

6.1 PERFORMANCE INDEXES AND BENCHMARKING METHODOLOGIES 140
6.1.1 Existing Standards and Comparison 141
6.1.2 Evaluation of energy-aware data plane 143
6.1.3 Evaluation of energy-aware control plane 145
6.2 THE ECONET TARGET ENERGY SAVINGS 148
6.2.1 The Adaptive Rate model 150
6.2.2 The Low Power Idle model 150
6.2.3 Using the optimal energy profile with AR and/or LPI primitives 151
6.2.4 Standby model 152
6.2.5 Considering the entire network infrastructure 152
6.3 ESTIMATING THE ENERGY SAVING GAINS 154
6.3.1 Overall energy consumption 155
6.4 SUMMARY 157

7 CONCLUSIONS 158

REFERENCES 160

1 Executive Summary

The energy consumption of the Internet is exploding as the number of devices connected to it increase in number, due to global coverage of an ever increasing population of users, as well as in consumption, due to the explosion of traffic volume fuelled by video and mobile data applications. Both the number and the capacity of the devices that make up the network depend on the peak load that the network has to carry during its peak hour. This peak capacity is directly connected with the peak energy consumption of the network. Looking however beyond the peak hour, it is clear that the load falls way beyond the peak, e.g., during local night-time. This drop of traffic however is not followed by a drop in the amount of energy consumed by the network. This is largely due to the lack of energy proportionality in the current generation of network devices that consume close to their peak energy independently of their actual traffic load.

Starting from the above observation, and projecting into the future demand for bandwidth based on recent reports, it is clearly shown that the energy consumption of the Internet will be one of the main challenges that technology will have to face in the future. Unfortunately, current and planned evolutions of the Internet, although they plan for a variety of improvements like increased capacity, easier manageability, and stronger security & privacy, do not address energy consumption issues. At the same time, strict requirements on energy efficiency are being pushed by policy makers and are sought by manufacturers and operators of devices and networks. Bridging the gap between the two is the main objective of the Energy COnsumption NETworks (ECONET) project whose mission, tools, and objectives are developed in detail in this first deliverable of WP2.

The first contribution of this deliverable towards achieving the said energy objectives is the characterization of the energy aware design space for networks and its breaking down into data plane related and control plane related actions. Within each category, a detailed account of energy aware techniques is provided, with further breakdown based on network layer, home, access, transport and core. At the data plain a combination of re-engineering, dynamic adaptation, sleeping, and energy efficient MAC design approaches are proposed, in order to meet the requirements set forth for network devices in different network parts. At the control plane, energy aware traffic engineering, resource allocation, and virtualization approaches are proposed for orchestrating together the energy conservation capabilities of individual devices and making the ensemble network be more energy proportional than its constituent elements.

The above techniques are integrated and positioned within three representative reference scenarios involving telecom operators and Internet service provider networks. The reference scenarios combine technologies in both planes and all layers and are thus suitable for a complete assessment of the energy-aware propositions of ECONET using a variety of performance indexes and benchmarking methodologies that are also detailed within the deliverable.

2 Introduction

2.1 Scope of the deliverable

This deliverable presents the user requirements as well as the technological background for the low Energy COnsumption NETworks (ECONET) that will be developed during the forthcoming work packages of the project. It thus creates the basic foundation for the development of green technologies for the data plane of network devices (WP3), the green abstraction layer for exposing them to higher level protocols (WP4), as well as the control plane strategies running higher up at the protocol level (WP5) for orchestrating the overall energy reduction efforts in ECONET. In addition, the deliverable defines the user requirements, the energy requirements, and the evaluation scenarios that will be presented to the integrated ECONET network and will be evaluated during the last stage of the project (WP6).

2.2 Objectives

The main objectives of the deliverable are three-fold:

Initially, it aims to demonstrate the importance of reducing the energy consumption in networks. This is achieved by analysing the constituent factors behind the current energy consumption of networks. First in line among these factors is the explosion in terms of the transmitted volumes due to the increase in the number of users in the network, but also, more important, in the traffic consumption of individuals. The latter owes largely to the rapid increase of video communications as well as to the explosive growth of traffic coming from (smart) mobile devices. The said traffic volumes translate to very high energy consumption, largely thanks to the energy disproportionality of existing network elements, which consume nearly the same amount of energy independently of whether they are completely idle or they carry the maximum amount of traffic that they can carry. To contrast the current dim situation of very high energy consumption, we present the desired energy consumption of future ECONET networks that will be developed during the project.

Secondly, the deliverable presents the available technological alternatives for achieving the energy consumption objectives. Achieving the energy reduction objectives set above requires taking drastic measures and solving non-trivial problems. The second objective of the project is to highlight the available options for dealing with the current high energy consumption of networks, both those developed within ECONET as well as elsewhere. We start by presenting existing proposals for the evolution of the Internet, most of which do not aim at energy consumption, but rather efficiency, manageability, and security/privacy objectives. These include optical backbone networks, content centric networks, autonomic networks, and virtualization. Next, we delve into energy targeting technologies at both the data and control plane of networks. At the data plane we look at the saving margins made possible by new technologies for the network terminations, the access, the switching, and routing network. In these levels, the proposed new technologies are based on re-engineering, dynamic adaptation (voltage scaling), sleeping and standby principles. Moving on to a higher layer, and looking at the control plane of ECONET, we discuss green traffic engineering strategies for routing traffic intelligently based on energy consumption objectives, as well as virtualization and efficient resource allocation approaches. Such approaches aim at aggregating traffic to a minimum amount of paths or using a minimum set of network resources (given QoS and fault tolerance objectives) and leaving the network elements in remaining paths enter sleep modes with the objective to save energy. Following the initial general treatment of energy saving technologies, we move on to describing the specific technologies and the targeted savings in which ECONET and its partners will focus their efforts.

Thirdly, the document defines the metrics and the reference scenarios for evaluating the results of the project. The last objective of the section is to define the metrics upon which the three different reference scenarios will be evaluated in the context of WP6.

2.3 Structure of the document

In addition to an Executive Summary (Section 1) and the Introduction (Section 2) the deliverable includes the following five sections:

Section 3, Energy Consumption in Next –Generation Networks: Trends and Technologies: it includes the overview of the existing Internet and the foreseeable evolutions in terms of technologies and traffic characteristics. As part of these evolutions it introduces for the first time the control and data plane techniques for energy conservation.

Section 4, Energy Efficiency Requirement for Next Generation Networks and Devices: the section presents a survey of energy consumption figures of existing devices at the different levels of the network hierarchy, from termination to core IP routers, as well as the requirements set for lowering the energy requirements of future devices. It also details the three reference scenarios of the project and connects them to requirements and technologies.

Section 5, Technology Specifications of Energy Efficient Devices: here the details of achieving energy reductions are presented. At the network device data plane the techniques including total re-engineering in order to cut down on power, dynamic adaptation of operating rate with response to current load, various sleeping techniques for devices that can go offline when the current load can be handled by a subset of them and energy efficient Ethernet, a new key technology appearing at various parts of the network. At the control plane, additional techniques include energy aware traffic engineering for selecting a subset of paths to carry the traffic, various joint resource (re)allocations around energy consumption and performance and virtualization techniques for adapting the function of network elements on the fly. The last part of the section gives concrete examples of the above with consumption profiles based on devices developed by ECONET partners which will be evaluated during the demos of WP6.

Section 6, Target Energy Savings and General Performance Indexes: this section puts together the metrics, requirements, and performance benchmarking methodologies that will be used throughout the project. Furthermore, it details an analytical framework for representing the impact of energy adaptive technologies and solutions for network devices under investigation by the ECONET consortium. Based on this framework, an evaluation is conducted for the potential impact estimated by the proposed model within the energy-efficiency degrees suggested by the ECONET consortium in the presence of the realistic scenarios for the access, home, metro/transport, and core networks.

Section 7, Conclusions: this section concludes the deliverable with a summary of the main technologies that will be examined within ECONET.

3 Energy consumption in next-generation networks: trends and technologies

In this section we first present a high level summary of the current state of the Internet based on the most recent published measurement reports which include the OECD Broadband Portal [55], the Akamai State of the Internet report [53], the Internet Inter-Domain Traffic report from the Atlas Observatory [54], and the Residential Broadband Internet Traffic report from T-Labs [56]. We discuss the evolution of Internet's topology (section 3.1), the exchanged traffic volumes, and traffic profiles (section 3.2). Moving on, we present data on the energy consumption of major Telecommunication companies as well as the objectives set for reducing it (section 3.3). The section concludes with the presentation of emerging network technologies including all optical, content centric, and autonomic networks (section 0), as well as a brief introduction of the design space for energy reduction in networks (section 3.5).

Identity of the employed reports:

- OECD: The OECD broadband portal provides access to a range of broadband-related statistics gathered by the OECD based on input from its members [55].
- AKAMAI: Akamai's globally distributed network of servers allows gathering massive amounts of information on many metrics, including connection speeds, attack traffic, and networking connectivity/availability/latency. The values of these data lies on that they represent the true quality of service that end users experience (as opposed to OECD data for example that gather only nominal capacities) [53].
- ATLAS: This report regards one of the first large scale longitudinal studies of Internet inter-domain traffic using direct instrumentation of peering routers across multiple providers. It has addressed significant experimental data collection and commercial privacy challenges to instrument 3,095 peering routers across 18 global carriers, 38 regional / tier-2, and 42 consumer and content providers in the Americas, Asia, and Europe. At its peak, the study monitored more than 12 terabits per second of offered load and a total of more than 200 exabytes of Internet traffic over the two-year life of the study (July 2007 to July 2009). Based on independent estimates of total Internet traffic volume it is believed that these probes directly monitor more than 25% of all Internet inter-domain traffic [54].
- TLABS: This report presents observations developed from passive packet-level monitoring of more than 20,000 residential DSL lines from a major European ISP. This unique vantage point provides a broad view of residential traffic, complementing the traffic characterizations based on backbone traces like [54] [56].

3.1 Internet topology evolution

The Internet has, in the last 3-4 years, seen a gradual transformation of its basic topological structure fuelled primarily from the success of content provider companies (Google, Yahoo!, Facebook, etc.), the prevalence of video and the content distribution networks that make its dissemination possible, and the rise of regional transit ISPs. The outcome of these transformations is a flattening of the -until recently- hierarchical structure, as well as a transformation to a mesh instead of a clear hierarchy.

Flattening of Internet's topology: Content providers like Google connect directly with access ISPs, thus most of the traffic flows directly to end customers instead of going through the traditional Internet core (the tier-1 providers). The 30 largest content providers generate around 30% of the observed inter domain traffic [54]. The result of this trend has been that the traditional hierarchical Internet (Figure 1) has been transforming into a more flat architecture (Figure 2).

Rise of video and CDNs: As will be elaborated next in the section, video, whether live or pre-stored, dominates -in terms of bytes sent- the traffic flowing through the Internet. This is largely made possible by Content Distribution Networks (CDNs), like Akamai or Limelight that maintain networks of servers around the world where local copies of content is kept and served to the local users. CDNs create a new intermediate layer between access and core tier-1 ISP. The addition of these networks complicates even more the topological structure of the Internet, making it more like a complicated mesh graph than a clear tree graph.

Rise of regional transit ISPs: A third fundamental change in the topology of the Internet is the rise of regional transit ISPs. This is an outcome of the consolidation in the telecommunications sector that has created in the last 10 years several intermediate / large players from the combination of once smaller access ISPs that consolidated through multiple acquisitions. These regional transit ISPs, like CDNs, eat out of the power and connectivity structure of the once dominant core tier-1 networks.

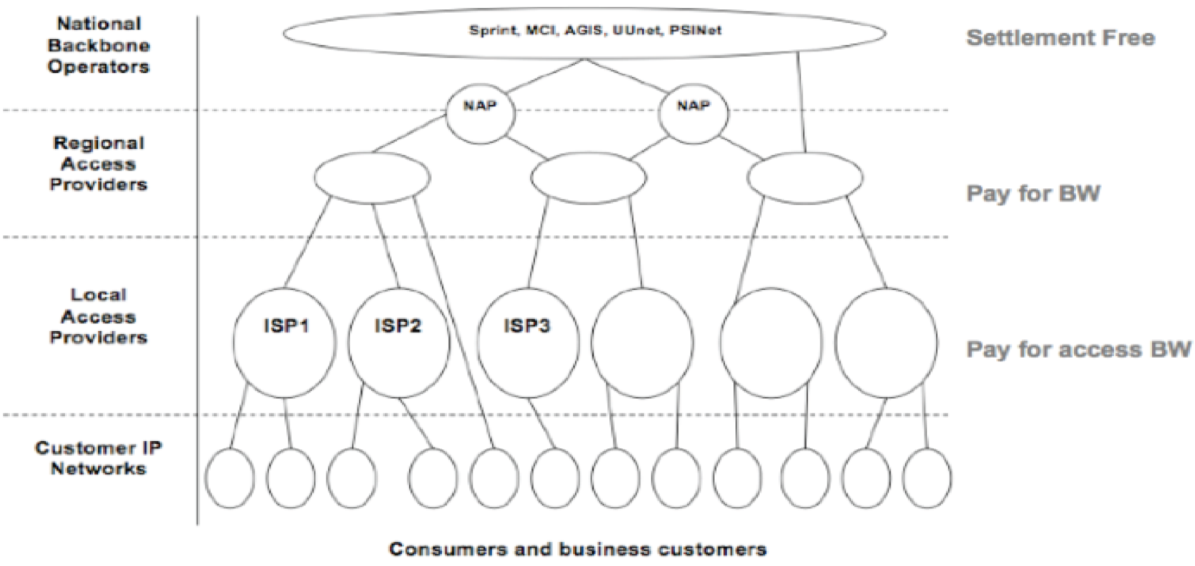

Figure 1: Traditional hierarchical Internet topology. Most traffic flowing between the core tier-1 providers.

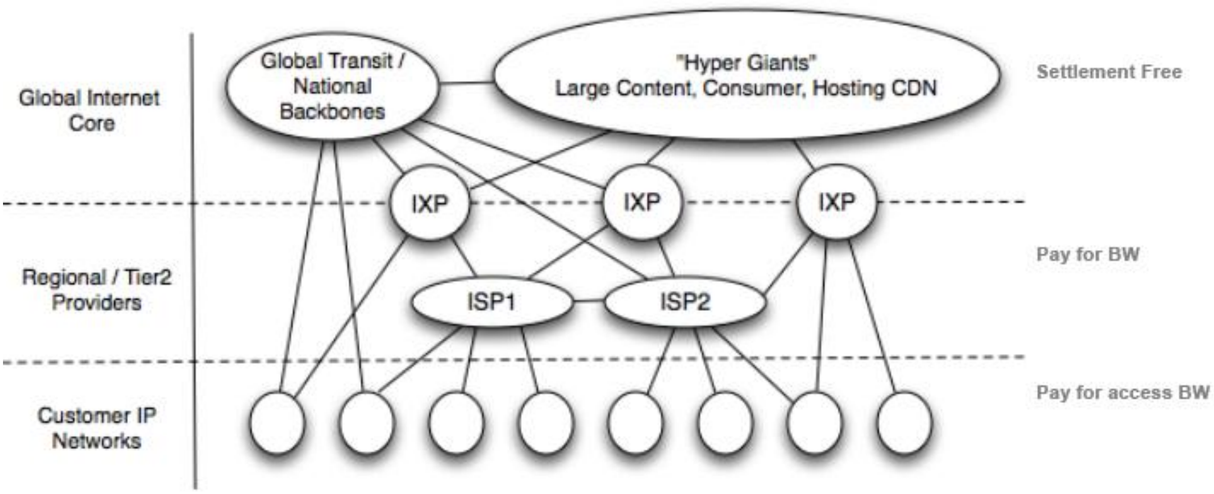

Figure 2: Emerging flat Internet topology. Most traffic flowing directly between content providers and customer IP networks while part of the traffic is flowing to the global Internet through Internet Exchange Points (IXP).

3.2 Traffic volumes and profiles

3.2.1 Traffic statistics based on protocol breakdown

Unlike traditional voice communication networks, the application and protocol mix of the Internet changes rapidly at very short time scales (even down to few months). This is a natural consequence of the openness of the platform that encourages innovation from the edge. For the same reason it is always challenging to come with predictions regarding the future traffic mix of the network. A first step however in this direction is to have a solid understanding of the current traffic mix. Next we summarize briefly the currently perceived picture regarding the application and protocol mix of the traffic that uses the Internet.

Most applications, whether they belong to traditional web applications, video applications or social networks, are consolidating over the HTTP protocol, making HTTP the largest contributor in terms of protocol traffic currently amounting to above 50% of all the traffic in the Internet (Figure 3). P2P protocols (like BitTorrent, eDonkey etc.) have increased their traffic in terms of absolute volume but percentage wise have shrinked to below 15% (Figure 3) whereas they amounted to 40-60% just a few years ago. This is a direct consequence of the very sudden rise of video traffic from web sites like YouTube.

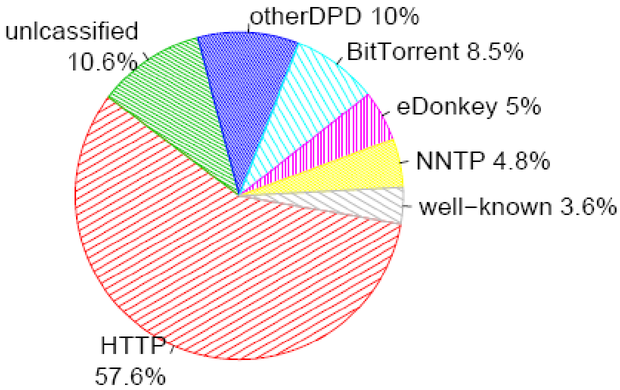

Figure 3: Protocol breakdown based on measured byte volumes [56].

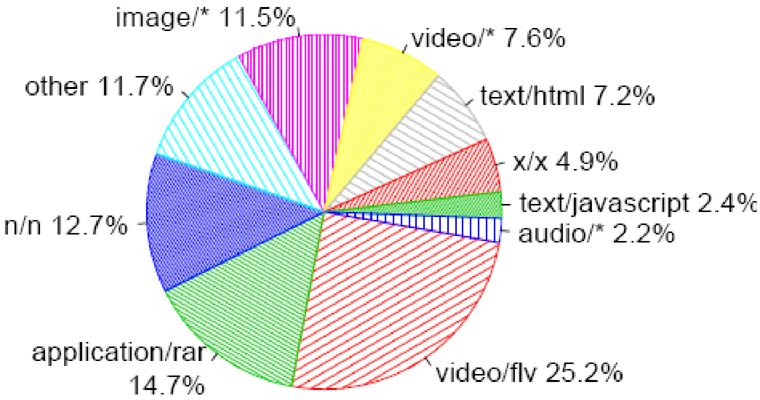

Figure 4: Content breakdown from analysing HTTP payloads [56].

Figure 4 presents the breakdown of HTTP traffic into different content/application types. The main point to notice is that Flash video (.flv) amounts to almost a quarter of all HTTP traffic. Another important contributor is the application/rar category with around 15% of HTTP traffic. Most of this traffic is due to HTTP downloads from One-Click-Hosting (OCH) sites like

RapidShare and MegaUploads. The recent success of such sites is a second reason for the decline of P2P's percentage in terms of overall traffic.

Figure 5 presents a table with the breakdown of origin sites for the observed HTTP traffic. The first four positions are occupied by video hosting sites like YouTube, followed by the Windows Update system of Microsoft and the CDN traffic carried by Akamai.

Rank	Domain	Fraction of Traffic
1	Direct Download Provider	15.3%
2	Video portal	6.1%
3	Video portal	3.3%
4	Video portal	3.2%
5	Software updates	3.0%
6	CDN	2.1%
7	Search engine	1.8%
8	Software company	1.7%
9	Web portal	1.3%
10	Video Portal	1.2%

Figure 5: Origin of traffic based on byte volumes [56].

3.2.2 Access speeds

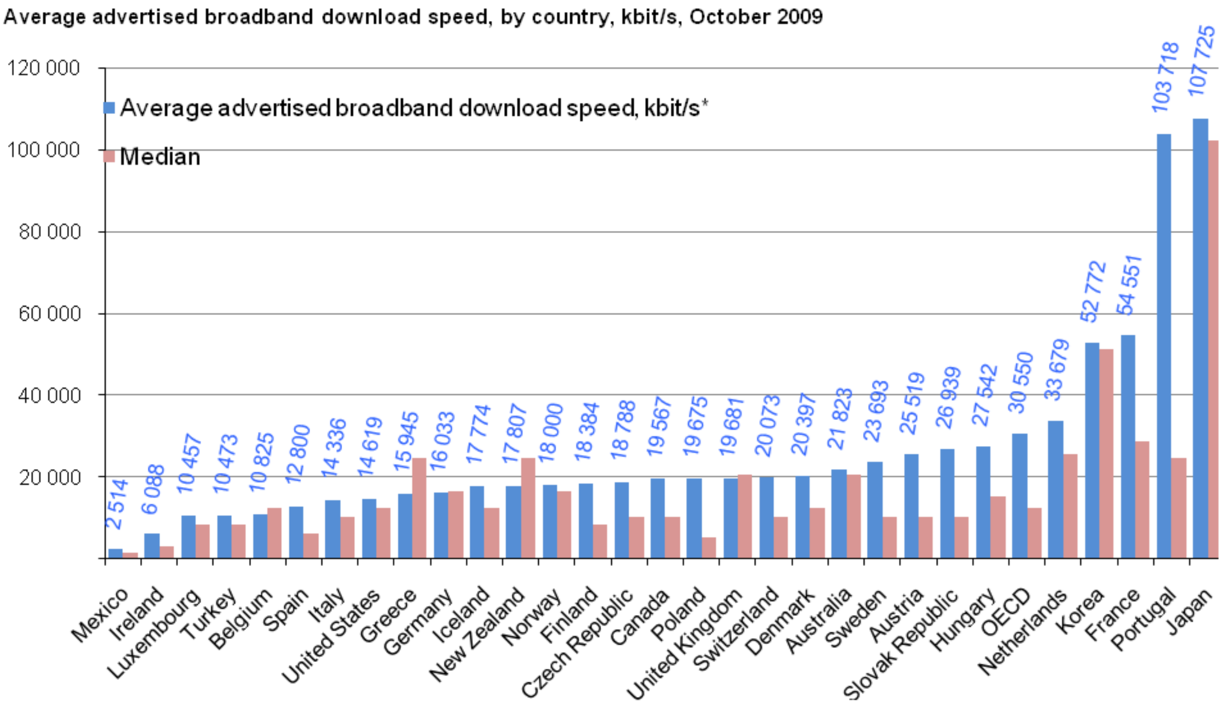

Figure 6: Advertised access speeds by country [55].

	Country/Region	Q2 '10 Avg. Mbps
–	Global	1.8
1	South Korea	17
2	Hong Kong	8.6
3	Japan	8.0
4	Romania	6.8
5	Netherlands	6.5
6	Latvia	6.3
7	Sweden	5.5
8	Czech Republic	5.3
9	Belgium	5.3
10	Denmark	5.2
…		
16	United States	4.6

	Country/Region	Q2 '10 Peak Mbps
–	Global	6.8
1	South Korea	38
2	Hong Kong	32
3	Japan	28
4	Romania	27
5	Latvia	20
6	Maldives	20
7	Belgium	19
8	Sweden	18
9	Monaco	18
10	Portugal	17
…		
11	United States	16

Figure 7: Average and peak measured download speeds by country [53].

Next we turn our attention to the access speeds at which residential customers connect to the network. We look at access speeds from two perspectives, their nominal capacities and the real measured speeds. The nominal speeds (capacity) are depicted in Figure 6 [55]. According to the collected data, Japan, Korea, Portugal and North European countries provide the fastest access speeds since they are the ones that have rolled out at large scale residential fibre connections.

Actual measured speeds from [53] regarding the average and peak download speeds for each country are depicted in Figure 7.

3.2.3 Traffic volumes and profiles

Regarding the generated traffic volumes, data is presented for inter-domain traffic and traffic from mobile users. Regarding inter-domain traffic, Google contributes 5% of observed inter-domain traffic according to [54] followed closely by Limelight and Akamai CDNs, Microsoft, and Carpathia hosting (MegaUploads, MegaVideo one click hosting sites). Traditional tier-1 ISPs have fallen to lower positions. The exact ranking is depicted in Figure 8. One of the major new findings of [53] is that over the last year the amount of traffic coming from Internet enabled mobile smart phones like the iPhone and the Android has been increasing very fast. This is a strong indication that in the future the traffic might be dominated by mobile users. It is important to notice that in several countries smart phones are already consuming on average several Gbytes of traffic per month.

In Figure 9, a typical graph is shown with the daily average of European and North American Internet traffic on a single daily timeline. In other words, 5am for Europe is 5am GMT and 5am for the US is 5am EDT. Both European and US traffic[1] is shown as a percentage of their respective peak traffic levels (i.e. 100% is the respective peak of each Europe and US traffic). As expected, both Europe and US Internet traffic have a lot in common. Both show regular, daily cyclical traffic patterns with Internet traffic dropping at night and growing during the day.

Rank	Provider	Percentage
1	Google	5.03
2	ISP A	1.78
3	LimeLight	1.52
4	Akamai	1.16
5	Microsoft	0.94
6	Carpathia Hosting	0.82
7	ISP G	0.77
8	LeaseWeb	0.74
9	ISP C	0.73
10	ISP B	0.70

Figure 8: Ranking of Internet autonomous systems according to the amount of observed inter-domain traffic [54].

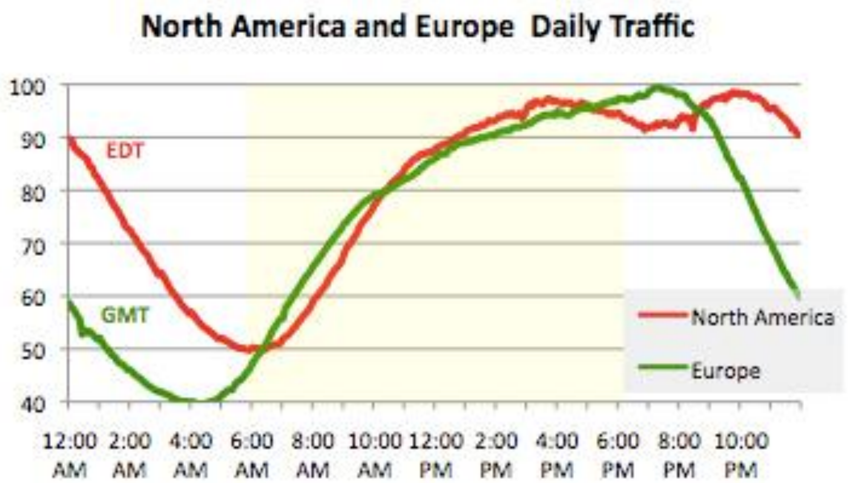

Figure 9: Daily average of European and North American Internet traffic.

3.2.4 Traffic projections for the network of tomorrow

Having summarized the current state of the Internet in the previous section, we move next to looking at predictions regarding the Internet of the future. Cisco has been publishing the Cisco Visual Networking Index [57], which has become the definitive source of information on the predicted traffic requirement over the coming years. The prediction is derived by drawing upon

[1] http://asert.arbornetworks.com/2009/08/what-europeans-do-at-night/

statistical data from the major market analysis firms and applying an elaborate forecast methodology that is detailed in [57]. Next we summarize the highlights of the forecasted traffic requirements as presented in [57].

Global Internet highlights

Global IP traffic has increased eightfold over the past 5 years, and will increase fourfold over the next 5 years. Overall, IP traffic will grow at a compound annual growth rate (CAGR) of 32 per cent from 2010 to 2015. Global IP networks will deliver 7.3 petabytes every 5 minutes in 2015. A growing amount of Internet traffic is originating with non-PC devices. In 2010, only 3 per cent of Internet traffic originated with non-PC devices, but by 2015 the non-PC share of Internet traffic will grow to 15 per cent. PC-originated traffic will grow at a CAGR of 33 per cent, while TVs, tablets, smartphones, and machine-to-machine (M2M) modules will have growth rates of 101 per cent, 216 per cent, 144 per cent, and 258 per cent, respectively.

Busy-hour traffic is growing more rapidly than average traffic. Busy-hour traffic will increase fivefold by 2015, while average traffic will increase fourfold. Global Internet video traffic surpassed global peer-to-peer (P2P) traffic in 2010, and by 2012 Internet video will account for over 50 per cent of consumer Internet traffic. As anticipated, as of 2010 P2P traffic is no longer the largest Internet traffic type, for the first time in 10 years. Internet video was 40 per cent of consumer Internet traffic in 2010 and will reach 50 per cent by year-end 2012. The sum of all forms of video (TV, video on demand [VoD], Internet, and P2P) will continue to be approximately 90 per cent of global consumer traffic by 2015.

Regional highlights

IP traffic in Western Europe will reach 19 exabytes per month by 2015, at a CAGR of 32 per cent. Monthly Internet traffic in Western Europe will generate 3.1 billion DVDs' worth of traffic, or 12 exabytes per month. IP traffic in Central and Eastern Europe will reach 3.7 exabytes per month by 2015, at a CAGR of 39 per cent. Monthly Internet traffic in Central and Eastern Europe will generate 0.8 billion DVDs' worth of traffic, or 3.1 exabytes per month.

3.3 Energy consumption and energy efficiency demand

It is estimated that the Information and Communication Technology (ICT) sector accounts for 2% of global carbon emissions [107]. 'Going green' is now a necessity for telecom operators, as energy expenses constitute nearly 25% of the total network costs. In addition to energy costs, restrictions on carbon footprints, and an increased emphasis on corporate responsibility are pushing Telco's to take measures towards energy consumption and carbon reduction. In fact, many telecom providers have started including such measures as a requirement in their requests for information and proposals from vendors. Decision-makers increasingly weigh the cost of powering a piece of equipment along with traditional features such as reliability, scalability and flexibility, resulting more and more often in earlier replacement of older equipment in favour of more energy-efficient devices. Thus, in the last few years, a large set of Telco's, ISPs and public organizations around the world reported statistics of network energy requirements and the related carbon footprint, showing an alarming and growing trend.

The overall energy consumption of Telecom Italia in 2006 has reached more than 2 TWh (about 1% of the total Italian energy demand), increasing by 7.95% with respect to 2005, and by 12.08% to 2004 [1]-[3]. In 2009, this consumption raised to 2.14 TWh. This energy consumption especially rose from network infrastructures, which contributed 70% of the total energy requirements. Datacentres weighed for 10%, while the remaining 20% is due to other spurious sources (e.g., offices, shops, etc.). Note that similar breakdowns of energy wasting can be certainly generalized for the largest part of telecom operators. Figure 10 reports the historical evolution of the energy consumption of TELIT, outlining how the energy requirements of telecommunication equipment are the major source of energy waste, and how, with the deployment of broadband access technologies, the consumption of customer premises are rapidly raising, becoming today a figure comparable to the direct energy consumption of telecoms.

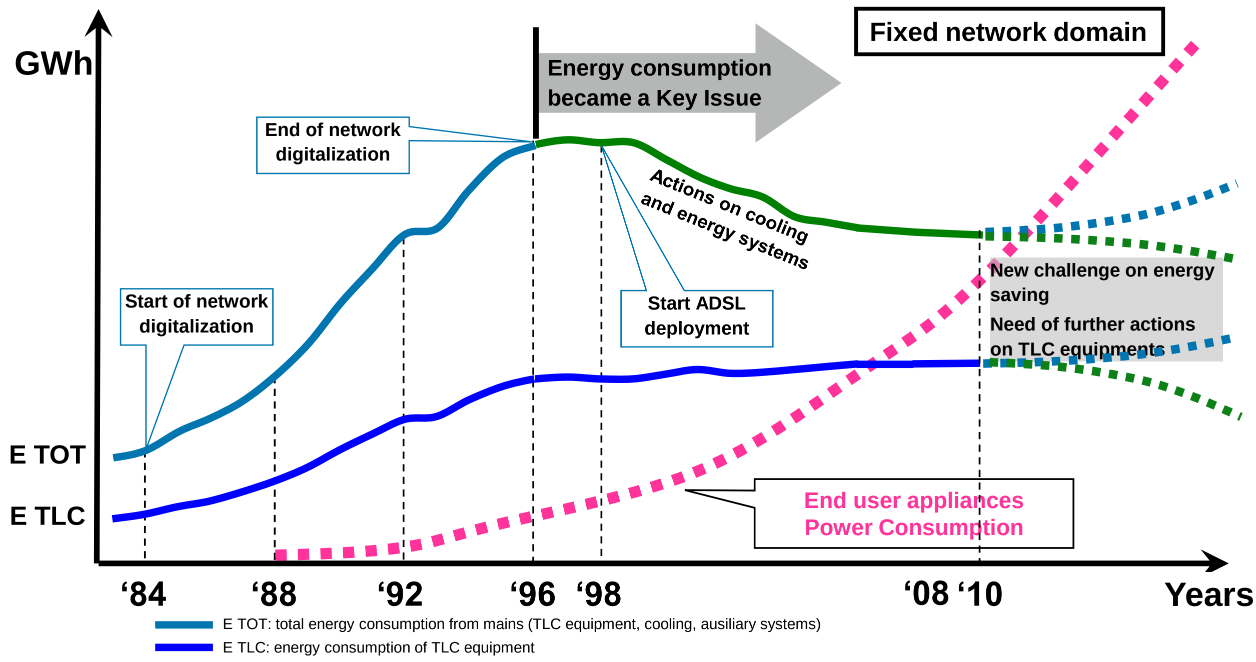

Figure 10: Electrical energy consumption evolution and future trends for TELIT's fixed network.

Another representative example comes from British Telecom, which reported energy requirements similar to the Telecom Italia ones: the overall power consumption for its network and estate during the 2010 financial year was 3.12 TWh [4] [5] (2,281 GWh of electricity to run UK networks, data centres and offices, 417 GWh of gas, heating oil and generator fuel at UK sites and 425 GWh of energy in countries outside UK). Moreover, it absorbed about 0.7% of the total UK's energy consumption in the winter of 2007, making it the biggest single power consumer in the nation [6]. About 10% of the UK's entire power consumption in 2007 was related to operating IT equipment [7].

The Deutsche Telecom group exhibits an overall energy consumption of 7.91 TWh during 2009. Only in Germany, Deutsche Telekom reported an overall amount of power consumption of about 3 TWh in 2007 [8], which increased of about 2% with respect to 2006 data. Deutsche Telekom justified this energy consumption increase as the result of technology developments (DSL), increasing transmission volumes and network expansion; though the figure also includes a small amount of spurious data, they outlined that almost 20% of such energy waste is due to cooling systems.

Moreover, the power consumption of Verizon during 2010 was 10.24 TWh, and during 2006 8.9 TWh (about 0.26% of USA energy requirements). AT&T weighs for 11.14 TWh in 2010, and declared to use 654 kilowatt hours (kWh) per terabyte of data carried on its network during 2008.

Table 1: Yearly energy consumption of some of the major Telecoms world-wide. These data were obtained by the Sustainability Report of each company.

Energy Consumption (TWh per year)					
Telecom	***2006***	***2007***	***2008***	***2009***	***2010***
Deutsche Telekom (World)	7.10	7.22	7.84	7.91	-
France Telecom (World)	3.66	3.47	4.57	4.38	-
Telecom Italia	2.10	2.15	2.13	2.14	-
British Telecom (UK)	1.94	1.99	2.03	2.28	2.28
British Telecom (World)	-	-	2.6	2.71	3.12
AT&T (World)	-	-	-	11.07	11.14
Verizon	8.90	-	-	10.27	10.24
NTT	-	2.76	2.76	2.75	-
Telefonica	1.42	-	4.76	5.05	6.37
SwissCom	-	-	0.43	0.40	0.40
China Mobile	-	-	9.35	10.62	11.94
SK Telecom	-	-	0.94	1.09	1.09

The requirements of France Telecom were about 4.38 TWh in 2009 [9], while it is committed to a 15% reduction in its global energy consumption by 2020 compared to the level in 2006. In 2010, the Telefonica group consumed 6.37 TWh, against a 2006 figure of 1.42 TWh, which amounts to 0.6% of Spain's total energy consumption [10]. The NTT group reports that the amount of electric power in fiscal year 2004 needed for telecommunications in Japan was 4.2 TWh [11], and their direct energy consumption 2.75 TWh in 2009. China Mobile energy requirements exceeded 10 TWh, raising of more than 1 TWh per year. Data regarding the yearly energy consumption of some of the major telecom operators world-wide are presented in Table 1.

In order to give further indications on these impressive figures, Table 2 reports the average price of the electricity according to the U.S. Energy Information Administration (EIA) agency for the years 2006-2010. These data clearly outline how the energy prices are increasing with a high pace, dramatically raising the request of efficiency from operators in the telecommunications field.

Table 2: Cost of the electrical energy according to the U.S. Energy Information Administration (EIA).

Average Energy Price					
	2006	***2007***	***2008***	***2009***	***2010***
Million $ per TWh	89	91	99	101	95

The figures above refer to the whole corporate consumption. As such, they account for numerous sources, other than the operational absorption of the networking equipment (e.g., offices' heating and lights). Nevertheless, we have included them to give an idea of the general trend.

Similar trends can be generalized to a large part of the other Telco's and service providers. The European Commission DG INFSO report in [12] estimated European Telco's and operators to have an overall network energy requirement equal to 14.2 TWh in 2005, which would have risen to 21.4 TWh in 2010, and to 35.8 TWh in 2020 if no green network technologies would be adopted. Thus, much likely as in other areas where energy efficiency is a concern, there are two main motivations that drive the quest for "green" networking:

1. the environmental one, which is related to the reduction of wastes, in order to impact on CO_2 emission (Figure 11);
2. the economic one, which stems from the reduction of costs sustained by the operators to keep the network up and running at the desired service level and their need to counterbalance ever-increasing cost of energy (Figures 12 and 13).

The Global e-Sustainability Initiative (GeSI) reported similar estimations [14], and weighed the carbon footprint of networks and related infrastructures at about 320 Mtons of CO_2 emissions in 2020. As shown in Figure 14, GeSI reported that, during 2002, network infrastructures for mobile communication and for wired narrowband access caused the most considerable greenhouse contributions, since each of them weighs for more than 40% upon the overall network carbon footprint. The 2020 estimation suggests that mobile communication infrastructures will represent more than 50% of network CO_2 emissions, while, as far as wired networks are concerned, both telcos' devices (e.g., routers, switches, etc.), and broadband access equipment will cause ever growing and non-negligible contributions, equal to 22% and to 15%, respectively.

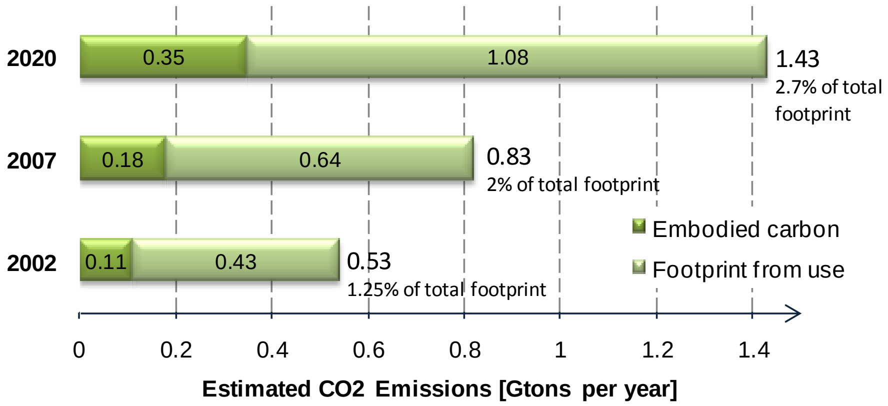

Figure 11: Estimate of the global carbon footprint of ICTs (including PCs, Telco's networks and devices, printers and datacentres). Source: Smart 2020 report by GeSI [14].

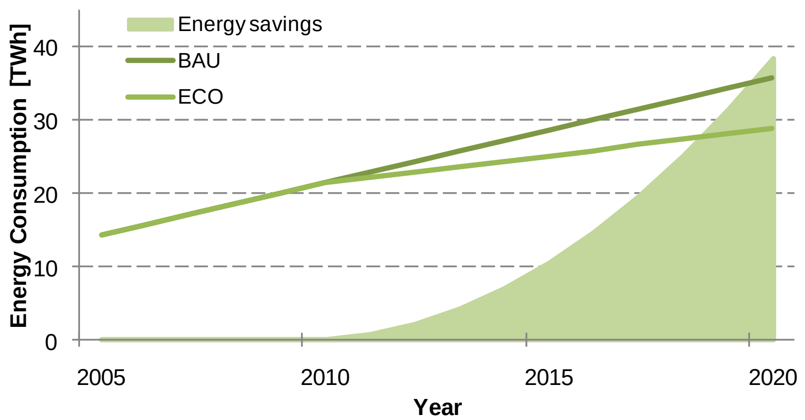

Figure 12: Energy consumption estimation for the European Telco's' network infrastructures in the "Business-As-Usual" (BAU) and in the Eco sustainable (ECO) scenarios, and cumulative energy savings between the two scenarios. Source: European Commission DG INFSO report in [12].

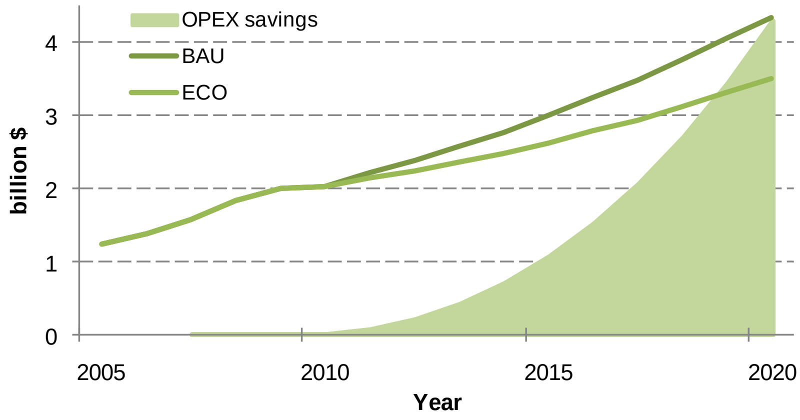

Figure 13: OPEX estimation related to energy costs for the European Telco's' network infrastructures in the "Business-As-Usual" (BAU) and in the Eco sustainable (ECO) scenarios, and cumulative savings between the two scenarios. Source: European Commission DG INFSO report in [12] and the U.S. EIA estimation on energy costs [13].

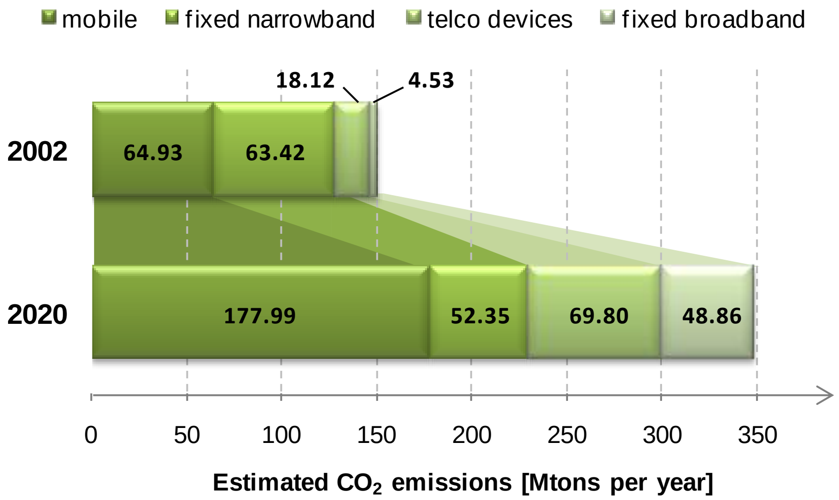

Figure 14: Greenhouse gas emission estimation according to GeSI [13].

Regarding wireless cellular networks, in 2008, researchers at NTT DoCoMo surveyed current mobile network energy consumption issues taking a Japanese mobile operator's network as an example [16]. The authors reported that the dominant part is due to radio access networks, and during 2006, each single user in NTT DoCoMo networks consumed, on average, 120 Wh per day for the network side and 0.83 Wh for the terminal side. The results of the investigation confirmed that current consumption can be reduced by introducing further IP-based Base Transceiver Stations (BTSs) and Radio-over-Fibre (RoF) technologies.

A deeper focus on energy consumption in today's and tomorrow's wired networks can be found in [17] and in [18], where Tucker *et al.* present a stimulating perspective on network design by focusing on cost and energy aspects. This perspective is based on a simple model for network technologies' evolution in the next years, and it uses the datasheets of state-of-the-art commercial devices, as well as projections of future broadband technologies.

By stating that today's network relies very strongly on electronics, despite the great progresses of optics in transmission and switching, the authors outlined how energy consumption of the network equipment is a key factor of growing importance. In this sense, they suggested that the ultimate capacity of the Internet might eventually be constrained by energy density limitations and associated heat dissipation considerations, rather than by the bandwidth of the physical components [19]. The authors pointed out that the data presented in their paper are based on a number of simplifications and approximations. Nevertheless, we believe that the results indicate some important trends.

The model aims at representing devices working at IP core, metro and access levels in order to evaluate the impact of different traffic grooming techniques both on scalability and costs, and on energy consumption of the overall network. To determine the energy consumption of the network, they have used information about the quantity of various types of networking equipment and the power consumption of these pieces of equipment, as minutely described in [20]. As outlined in Figure 15, the model exploitation highlighted that, not surprisingly, the overall energy consumption will increase as the capacity of the network expands. In this respect, it is worth noting that today's average access rates are about 2 Mbps. Thus, starting from the data in Figure 15, today energy requirements of access networks account twice with respect to the core.

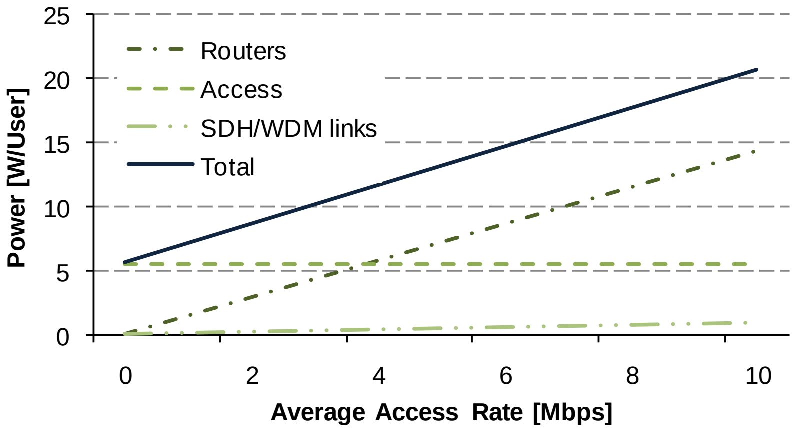

Figure 15: Average power consumption per user with respect to the increase in network access rate according to the results in [21].

In today's Broadband access network, the energy consumption is dominated by the energy in the user modem. The Passive Optical Network (PON) provides the lowest energy solution for broadband access with respect to point-to-point Ethernet, Fibre-To-The-Node (FTTN) and Wi-Max technologies. This feature of PON could become a driver of future PON deployment in response to concerns about the greenhouse impact of the Internet. Moreover, the authors have shown elsewhere [21] that at low access rates the power consumption in DSL networks is similar to that of PONs.

Tucker *et al.* finally demonstrated that the energy consumption in the routers — particularly in the core — will become more significant, as user access rates increase. The energy consumed in WDM links is relatively small. There is little evidence that optical burst switching or optical packet switching will significantly reduce the cost or energy consumption in future high capacity networks.

In this respect, Tucker's conclusions are confirmed by the trends shown in Figure 16, which reports that capacities and power consumptions of high-end routers grow in an exponential way, by a factor of 2.5 and 1.65 every 18 months respectively, on a per rack basis. In more detail, the data in Figure 16 has been obtained by [27] and completed with the datasheets of recent commercial top routers (e.g., Cisco CRS-1 [22], Juniper T1600 [23], Huawei Quidway 5000E [24], Brocade NetIron XMR 160000 and 320000 [25], etc.).

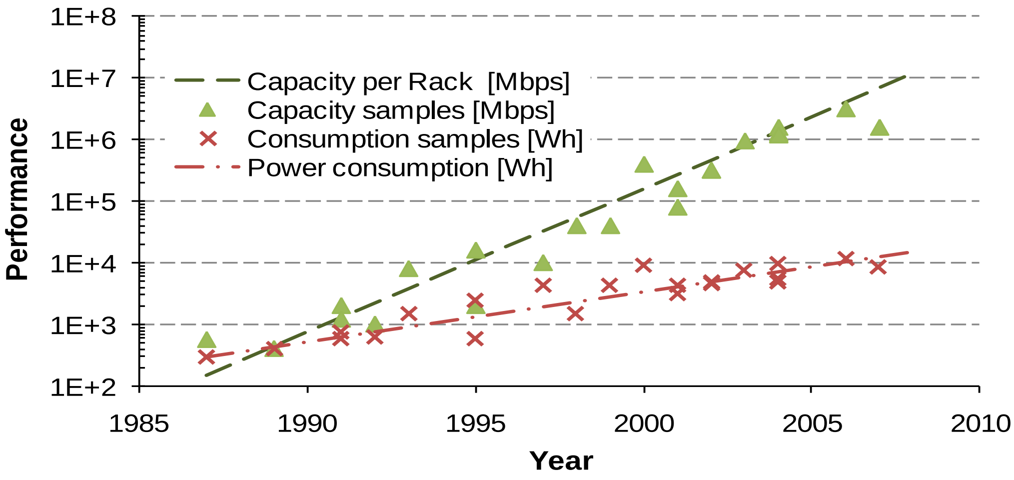

Figure 16: Evolution of capacities and energy requirements of high-end routers from 1985 to 2010. The estimation is on a per-rack basis [22] [23] [24] [25].

To support new generation network infrastructures and related services for a rapidly growing customer population, Telco's and ISPs need an ever-larger number of devices, with sophisticated architectures able to perform increasingly complex operations in a scalable way.

As just noted, high-end IP routers are even more based on complex multi-rack architectures, which provide more and more network functionalities and continue to increase their capacities with a factor of 2.5 every 18 months [1]. At the same time, as shown in Figure 17 and as suggested by Dennard's scaling law [2], silicon technologies (e.g., CMOS) improve their energy efficiency with a lower rate with respect to routers' capacities and traffic volumes, by increasing with a factor of 1.65 every 18 months.

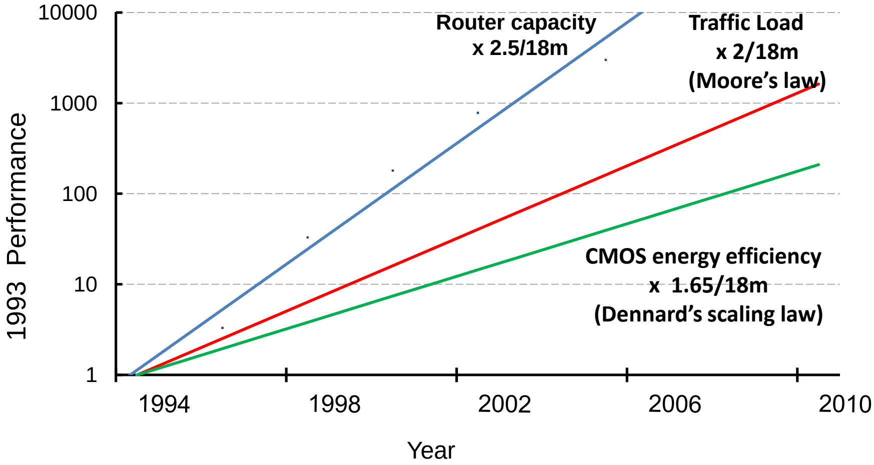

Figure 17: Evolution from 1993 to 2010 of high-end IP routers' capacity (per rack) vs. traffic volumes (Moore's law) and energy efficiency in silicon technologies [1].

3.4 Next generation networks technologies

This subsection regards a short description of emerging technologies that will dominate in next generation networks and are able to enhance the energy efficiency of these networks. These technologies include the transition of the core networks of ISPs and telecom operators to optical backbone networks, the incorporation of autonomic characteristics, and the design of self-* functionalities for optimization purposes from the energy consumption perspective and the trend towards the establishment of a new way of communication in the future networks based on the content generated or requested from a user.

3.4.1 Optical backbone networks

Contemporary backbone data networks are mainly using distributed IP routing so as to move packets from source to destination. Moreover, the transport network is typically used for providing static point-to-point high capacity connectivity among IP routers. This approach has been in place for many years and its main advantage is that it draws a concrete line among the IP and transport worlds, which are typically managed and configured by different Service Provider teams. From the energy consumption point of view, this backbone architecture is far from being optimal since it is heavily based on power hungry routers. Furthermore, as the Internet continues to grow, it requires network elements with larger capacity, higher transmission rate and faster processing speed.

With the exponential traffic increase in IP traffic, remaining at lower layers when possible is advantageous for operators wanting to keep their energy bills under control. As a result, the most eco-efficient architecture is a multilayer one that can automatically direct traffic to the lowest level of switching required according to bandwidth, network availability and service requirements (Figure 18). A network design approach that is currently promoted, especially by transport equipment vendors, in order to address contemporary backbone data networks power consumption inefficiency, is called "router by-pass" or "router off-load". The main idea of this approach is that transport equipment's intelligence is enhanced so as to be able to dynamically establish high capacity circuits minimizing the number of intervening IP routers.

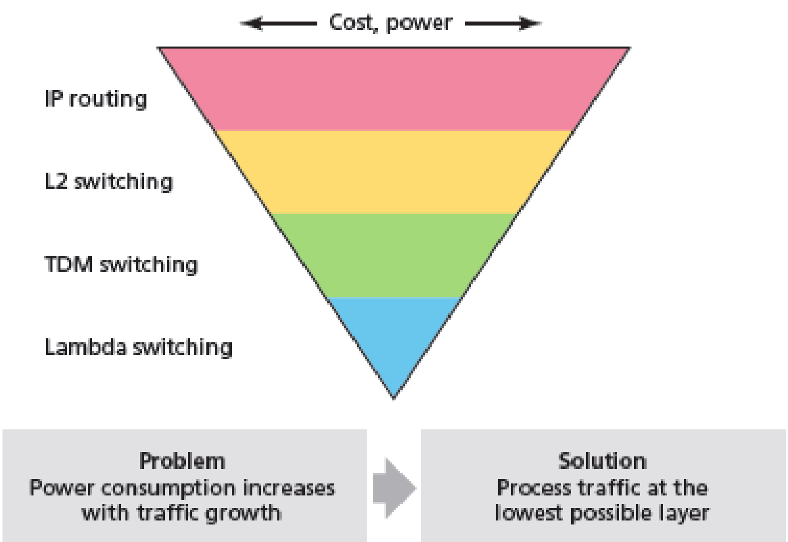

Figure 18: Power consumption through the transport and switching layers.

Towards this direction, optical backbone networks constitute the best candidate since they can exploit the enormous bandwidth that may be provided by fibre optical technologies (order of Tbps) and the flexibility for establishment of end-to-end optical circuits among various nodes in the network. Wavelength-division multiplexing (WDM) is a technology that is being used for multiplexing a number of optical carrier signals onto a single optical fibre by using different wavelengths (colours) of laser light. This technique enables bidirectional communications over one strand of fibre, as well as multiplication of capacity. Thus, the creation of optical paths within optical backbone networks has been utilized for the dynamic establishment of high capacity circuits with reduced energy demands (Figure 19) [58] [59] [64].

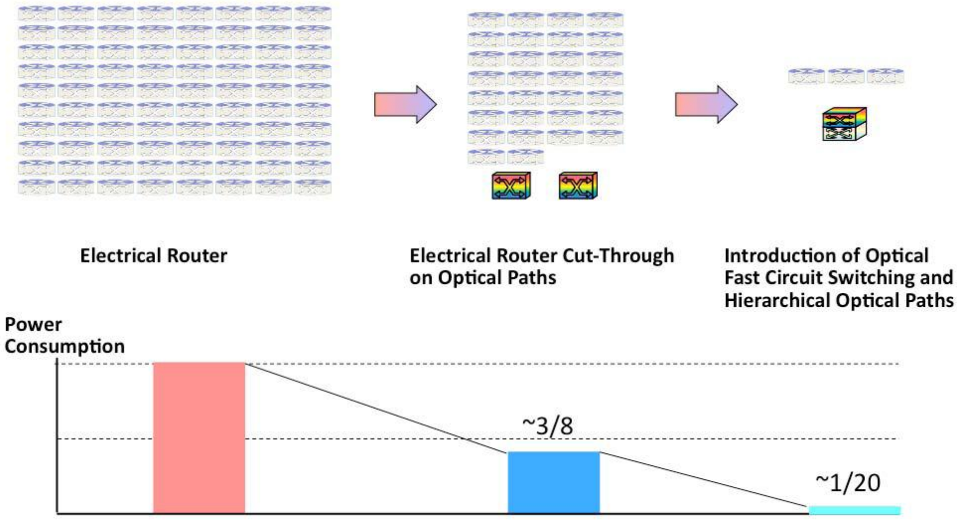

Figure 19: Electrical router cut-thorough and optical fast circuit switching [58].

Optical Transport Networks (OTN) also provide efficient sub-wavelength bandwidth management capabilities. Multiple transport options are available for individual management of traffic relations generated in the IP routing layer, which are typically fractions of 10G, 40G or future 100G DWDM line rates and need individual forwarding, operations, monitoring and SLA (Service Level Agreement) assurance according to their service mixes and destinations. OTN's features provide a transport foundation for IP traffic relations on which router ports and even sub-ports can be mapped to the most optimal transport entity: a wavelength (optical channel), a fixed-rate virtual container (optical data unit, ODU) or a variable-rate virtual container (ODUflex).

By augmenting such architectures with GMPLS technology that enables control plane integration, operators can automate the selection of the most power-efficient layer. While GMPLS today is used only at the electrical level, Alcatel-Lucent extends it to the photonic layer, leveraging its resilience features and capacity for resource optimization. This directly contributes to lower power consumption.

Regarding the network control plane, it may also be utilized efficiently for reducing energy consumption. An ASON/GMPLS (automatic switched optical network/generalized multi-protocol label switching) optical control plane simplifies network operations with the goal of creating a 'self-running' network in which 'the network is the database'. With ASON/GMPLS, the network has the intelligence to choose the most power-efficient layer for transport. Furthermore, with GMPLS provisioning and restoration capabilities at the photonic level, operators can improve their SLA performance and the quality of their wavelength services. GMPLS control plane intelligence enables dynamic service provisioning and improves bandwidth monetization through better utilization of network resources. In essence, GMPLS provides the operator with a photonic network that possesses the flexibility, automated operations and resilience typical of digital based networks — along with the ability to forward traffic at the lowest cost per bit without significant limitations.

By enabling resilient, automated and power-efficient networks, GMPLS brings a number of CAPEX and OPEX advantages in addition to eco-benefits. Specifically, it provides for:

- Lower power consumption due to greater cross-layer intelligence for resource optimization.
- Delegation of several OSS processes to the control plane for automation including discovery processes for network topology, resources and services, end-to-end connection routing for optimal resource utilization, flow-through service provisioning, and mesh restoration.
- An intelligent restoration mechanism that boosts network reliability and allows network failures and fibre cuts to be accumulated and fixed in batches, instead of one at a time as happens today. This allowance for planned network maintenance activities helps reducing the cost of on-site maintenance as well as travel-related CO_2 emissions.
- Fewer site visits for provisioning.
- A smaller footprint due to the use of fewer network elements.

A variety of new component technologies is also designed that may be integrated into optical platforms and combine multiple protocols into each individual device. Specific component technologies include lower power cooling fan units, power-efficient DC/DC converters and chips, lower power optical components and dynamic power and thermal management technologies. Innovative photonic OAM features (e.g. Zero-Touch Photonics) are also introduced. The Zero Touch Photonics (ZTP) is a new concept proposed by Alcatel-Lucent. It consists of OAM features for complete networking capabilities at the photonic layer without requiring on-site intervention. The chief characteristics of ZTP include photonic switching (T-ROADM and ROADM architectures: colourless, directionless, multi-degree), photonic OAM with fault localization and performance monitoring, photonic design tools for end-to-end life cycle management, photonic restoration via GMPLS and adaptable optical transport.

3.4.2 Next generation technologies in the access part

Access networks are typically used for connecting end users to the Internet gateway. An access network can be a fixed line network, in case a carrier cable reaches the end user, or a wireless network.

The most popular fixed access technology is the copper wire based Digital Subscriber Line (DSL) along with its variants Asymmetric DSL (ADSL) and Very-high-bitrate DSL (VDSL). ADSL manages to achieve up to 24Mbps downstream speeds and 1Mbps upstream speeds over a range of 5.5 km. VDSL achieves increased capacity, by utilising a wider part of the spectrum reaching up to 52 Mbps downstream and up to 16 Mbps upstream, but over shorter distances -around 300 m. Another fixed access technology that is constantly increasing worldwide commercial penetration is Passive Optical Networks (PONs); it is based on fibre optic technologies and passive optical couplers/splitters so as to reduce the overall length of the deployed fibre and keep the overall network cost low. As a result, a section of the deployed access fibre network is shared among the PON users. The PON basic idea is illustrated in Figure 20.

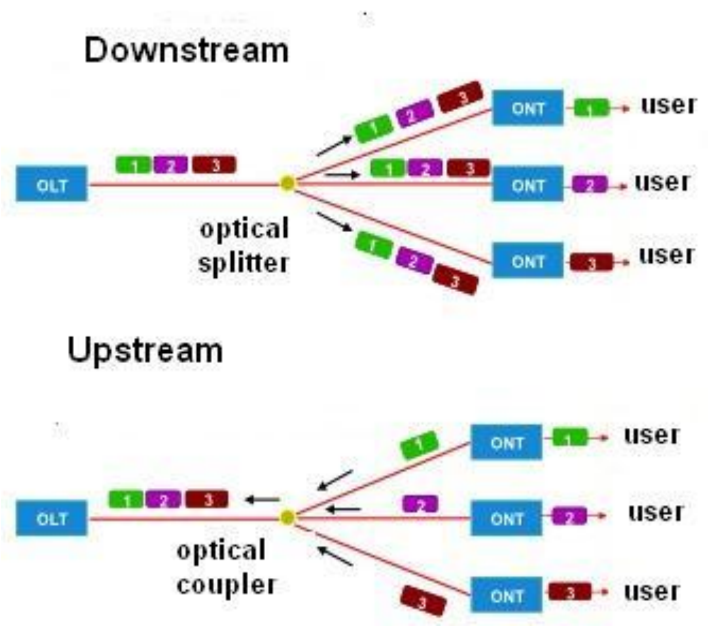

Figure 20: PON operation principle.

PON rates can scale up to 10Gbps and can be shared among the PON users over a range of 10-20km. Typically, up to 64 users can co-exist on a PON. The most deployed PON technological alternatives are GPON and EPON. PON technology can be combined with DSL, acting as the DSL backhaul; in this case the optical network does not reach the end user premises, but either the local cabinet (Fibre To The Cabinet – FTTC) or the local building (Fibre To The Building – FTTB). This is illustrated in Figure 21 and refers to Telecom Italia's FTTx penetration plans as presented in [1].

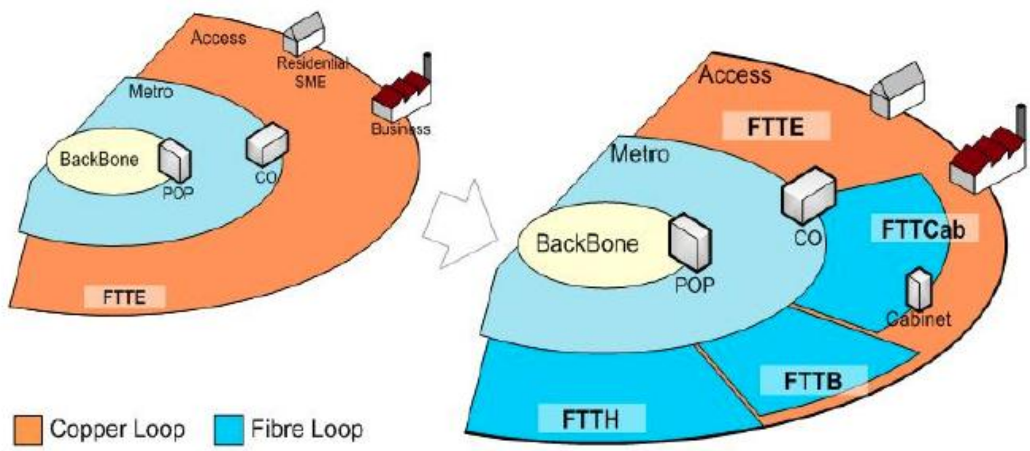

Figure 21: Access network topology evolution in Telecom Italia [1].

Technology	Range (km)	Bit rate (Mb/s)	Users/node	Minimal user density (subs/km^2)	Power/subs (with PUE) (W/subs)
ADSL ADSL2+	5.5 1.5	8[1] 24[1]	384–768	4–8 50–100	2–4
VDSL VDSL2+	1.0 0.3 0.3	26[1] 55[1] 100	16–192	5–60 50–700 50–700	6–10
GPON (32) GPON (64)	20 10	2488/32 2488/64	(4–72) * 32 (4–72) * 64	0.1–2 0.8–14	0.4–1.6
Mobile WiMAX HSPA LTE	0.340 (3 Mb/s) 0.240 (3 Mb/s) 0.470 (3 Mb/s)	1–70 1–14 1–300	272[2] 225[2] 180[2]	N/A N/A N/A	27[3] 68[3] 18[3]

[1] Downstream [2] Simultaneous Active Users [3] Modelled for 300 subscribers per km^2

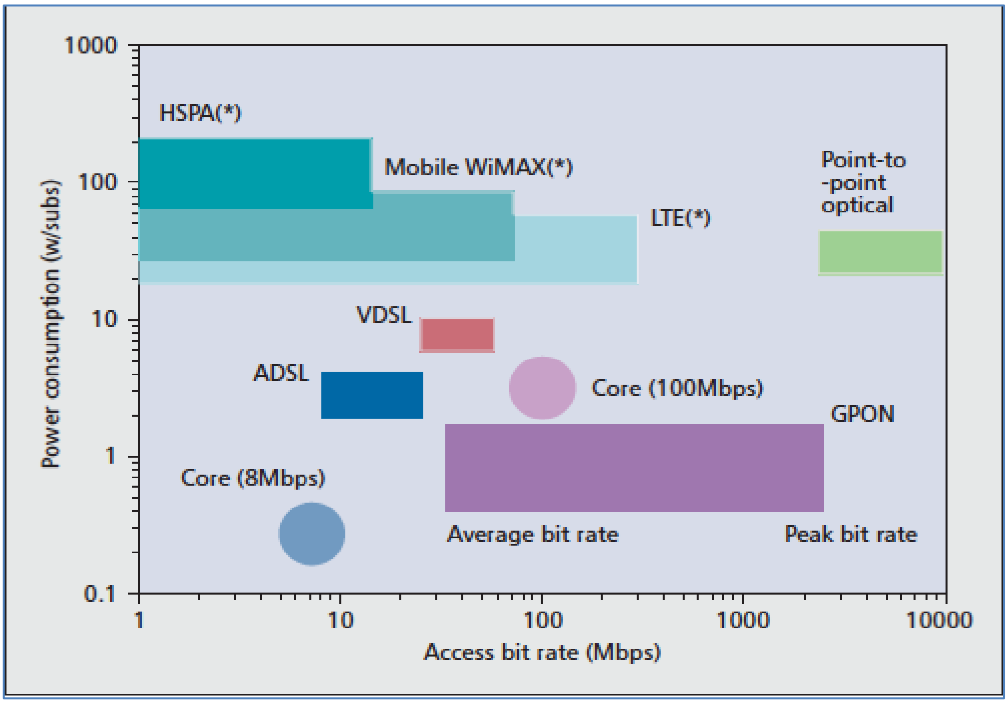

Figure 22: Access technologies power consumption [100].

Key technologies for wireless access networks include Worldwide Interoperability for Microwave Access (WiMAX), High Speed Packet Access (HSPA) and Long Term Evolution (LTE). WiMAX is based on the IEEE 802.16 standard and can support rates up to 70Mbps to mobile users. HSPA is the successor of UMTS and can support peak rates up to 14 Mbps in the downlink and 5 Mbps in the uplink. LTE, which is considered to be the fourth generation radio technology, can theoretically achieve up to 326 Mbps rates in the downlink and 86 Mbps in the uplink.

Figure 22 illustrates power consumption values for fixed and wireless access technologies as presented in [100]; they include the overhead power that needs to be consumed in order to guarantee proper equipment operation (e.g. cooling). The authors assumed that the overhead power is the same as the one needed for network equipment to operate. Moreover, in terms of power consumption at wireless access networks, the majority of power is consumed at the base station and is highly depended on the local mobile users' density. Figure 22 presents base station consumptions assuming operation at an urban-suburban area with user density 100-300 users/km^2. Note that the order of power consumption among wireless technologies can change if the assumption for user densities is modified.

3.4.3 Content centric networks

Content-centric networking (CCN) is an alternative and innovative approach to the architecture of computer networks. In CCN, the principal paradigm is not end-to-end communication between hosts as it is in the current Internet architecture. Its founding principle is that a communication network should allow a user focusing on the data he or she needs, rather than having to reference a specific, physical location where that data is to be retrieved from. This stems from the fact that the vast majority of current Internet usage (a "high 90% level of traffic") consists of data being disseminated from a source to a number of users.

CCN replaces host-to-host conversations with named data oriented communications by introducing content routers into the network. While the basic operation of a content router is very similar to that of an IP router, the key departure is that it supports name-based routing and caching for content retrieval throughout the network. This way, content-centric networking obviates the need of deploying pre-planned, application specific mechanisms such as Content Delivery Networks (CDNs) and P2P networks, which require sophisticated network services for mapping named content to hosts. Furthermore, this approach can reduce transit traffic in the backbone network links since the establishment of host-to-host communication for dissemination of content is not necessary. Thus, the performance of content centric networks in terms of energy efficiency as well as its impact in network based QoS parameters has to be further studied and elaborated.

CCN uses user-friendly, hierarchical names like URL. The basic operation of a CCN router is very similar to an IP router. An interest (or request) packet arrives at an interface, and then a look-up is done on its name to make a forwarding decision (request from node H to node A in Figure 23). Since CCN routers support caching within the network, any intermediate CCN routers on paths that have the requested chunk can answer the request (e.g. nodes B and D in Figure 23). In case of a new request (request from node I to node A in Figure 23) the content can be provided by intermediate routers (node D in Figure 23).

In [65], it is shown that a change in network architecture from host-oriented to content-centric networking (CCN) can open up new possibilities for energy-efficient content dissemination. However, CCN's mechanisms need to be efficiently implemented to keep the energy consumption for content processing as low as possible, while a sufficient number of content routers must be deployed throughout the network to reap the benefit of content caching. The results confirm that CCN is more energy efficient than conventional CDNs and P2P networks, even under incremental deployment of CCN-enabled routers.

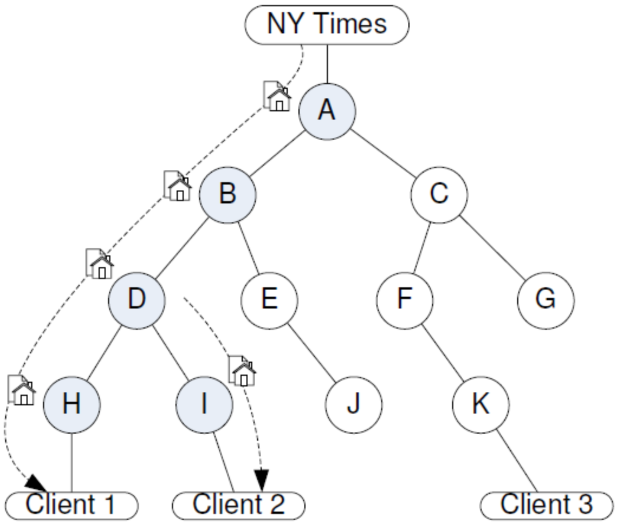

Figure 23: Routing in content centric networks [65].

3.4.4 Autonomic networking and Cloud computing

Autonomic networking is able to offer a basic platform to achieve greener Internet services, since it looks for better resource management and utilization according to the status of the system. In autonomic networks, network components may interact with each other in order to find optimal ways of collaboration from the energy perspective. Through the design of self-* functionalities (self-optimisation, self-healing, self-awareness), the network nodes may collaborate, exchange information regarding their status and take distributed decisions. These decisions can be based on the achievement of a, local or network level, business or technical goal (e.g. in the case of green networks the achievement can be the reduction of total energy consumption without significant implications to the network performance). At the local level, each node is able to monitor its status and interfaces and self-optimise its functionality. At the network level, nodes have to collaborate and act as parts of a global team aiming at total energy consumption reduction.

In addition to autonomic networking concepts, cloud computing is also considered nowadays as a key component of next generation networks. Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction [65]. Cloud computing contributes to reduction of energy consumption for processing and storing data, since part of computing and storage resources may be highly utilized while the remaining part can be in sleeping status. Techniques like sleep scheduling and virtualization of computing resources improve the energy efficiency of cloud computing. However, even with large savings in energy consumption that may be achieved through the realization of storage and processing in data centres, energy consumption in the transport and switching part of the network may be increased in parallel, due to the increase in the transferred amount of end-to-end traffic [67].

Taking these considerations into account, it is claimed in [67] that cloud computing may be proven energy efficient for the provision of specific services. It is shown that cloud computing can enable more energy-efficient use of computing power, especially when the users' predominant computing tasks are of low intensity or infrequent. However, under some circumstances, cloud computing can consume more energy than conventional computing where users perform all computing on their own PCs. Thus, even with energy-saving techniques such as server virtualization and advanced cooling systems, cloud computing cannot be considered always as the greenest computing technology. A summary of conditions under which energy consumption is significant in transport, storage, and processing for the provision of public and private cloud services is provided in Figure 24.

Energy Component	Service type	Software as a Service	Storage as a Service	Processing as a Service
Transport	Public	High frame rates	Always	Medium to high encodings per week
	Private	Never	High download rates	Never
Storage	Public	Never	Low download rates	-
	Private	Never	Low download rates	-
Processing	Public	Few users per server	Never	Medium to high encodings per week
	Private	Few users per server	High download rates	Medium to high encodings per week

Figure 24: Conditions under which energy consumption is significant [67].

3.5 Energy-aware design space

The problem of improving a network's energy efficiency requires the synergy of several technologies under the consideration of energy consumption as the primary objective. A comprehensive treatment of this multidimensional problem is necessary in order to be able to provide a solution that offers considerable saving [69].

The most crucial energy-related improvements in network equipment design mainly refer to the data-plane of network equipment since this usually includes the most energy starving hardware elements. **Energy-aware design of the network device data plane functions** is clearly one of the most important areas of the energy aware design space. On a second level, exploiting the equipment's novel green capabilities/functions, as well as the overprovisioned design of network infrastructures to develop local and network wide control strategies in order to minimize overall energy consumption while meeting the operational needs, forms another important design area termed as **Energy-aware strategies design at the control plane**. The development of control strategies should be clearly facilitated by an abstraction of the green capabilities/functions of heterogeneous equipment exposed through an interface referred as the **Green abstraction layer**.

In the following, the two areas of energy-aware design recognized as the most crucial, energy-aware design of the network device data plane functions and energy-aware strategies design at the control plane, are further explained.

3.5.1 Energy-aware design of the network device data plane functions

According to [70], at the highest level, energy-related improvements in network equipment design can be classified as *organic* and *engineered*. Organic efficiency improvements are commensurate with Dennard's scaling law; every new generation of network silicon packs more performance in a smaller energy budget. Engineered improvements refer to active energy management including -but not limited to- idle state logic, gate count optimization, memory access algorithms, I/O buffer reduction, and so forth.

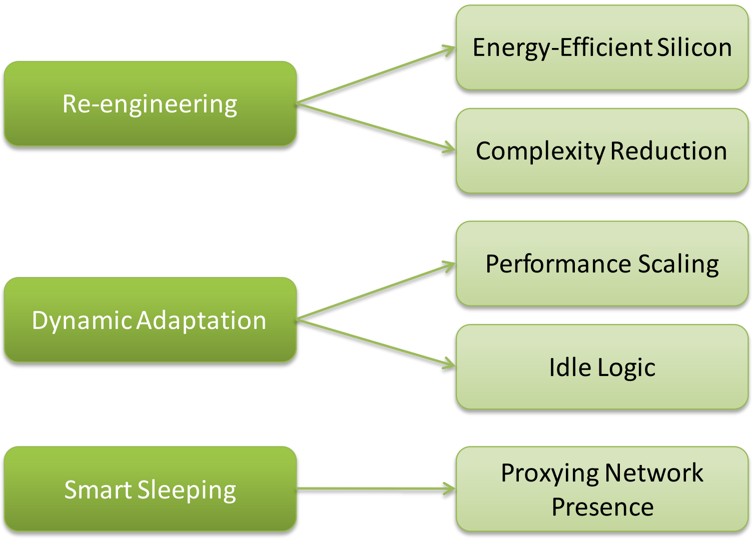

Figure 25: Taxonomy of undertaken engineered approaches for the energy efficiency of the Future Internet devices.

In this framework, the largest part of undertaken approaches regarding engineered improvements is funded on few base concepts, which have been generally inspired by energy-saving mechanisms and power management criteria that are already partially available in computing systems. These base concepts for energy efficiency in wire-line networks can be classified as shown in Figure 25 into Re-engineering, Dynamic adaptation and Sleeping/standby approaches.

Re-engineering approaches aim at introducing and designing more energy-efficient elements for network device architectures, at suitably dimensioning and optimizing the internal organization of devices, as well as at reducing their intrinsic complexity levels.

The *dynamic adaptation* of network/device resources is designed to modulate capacities of packet processing engines and of network interfaces, to meet actual traffic loads and requirements. This can be performed by using two power-aware capabilities, namely, *performance scaling (dynamic voltage and frequency scaling)* and *idle logic*, which both allow the dynamic trade-off between packet service performance and power consumption.

Finally, *Sleeping/standby* approaches are used to smartly and selectively drive unused network/device portions to low standby modes, and to wake them up only if necessary. However, since today's networks and related services and applications are designed to be continuously and always available, standby modes have to be explicitly supported with special proxying techniques able to maintain the "network presence" of sleeping nodes/components and to guarantee a short "wake-up" time.

These approaches are not exclusive among themselves, and it is considered that research efforts will be needed in all such directions in order to effectively develop new-generation green devices.

3.5.2 Energy-aware strategies design at the control plane

Currently, network design is based on deploying and maintaining densely interconnected infrastructures built for reliability and high performance. Typically, high-end routers are used in the core, lower-end distribution routers around the core and even lower-end access routers and switches at the periphery. Understanding the power demands of current networking equipment under different configurations and traffic loads provides the opportunity of replacing power-hungry systems with lower power systems that still provide required reliability and performance.

Towards this direction, virtualization technology [79] is considered as a powerful tool. Virtualization refers to the division of a hardware platform into several isolated virtual environments. The virtualization capabilities of the network equipment [93] such as routers and switches are the basis for developing network consolidation strategies achieving this way more efficient utilization of resources while permitting operating less energy consuming physical equipment.

Designing control strategies from the energy perspective is the final step towards improving a network's energy efficiency, since the illustrated network and the power-aware capabilities of the network components are taken into account to devise mechanisms for achieving energy savings. Energy aware traffic engineering approaches [68], as an extension of traditional traffic engineering, fall into this category, since they are targeted to effectively serve traffic while reducing the network's overall energy footprint. The characteristics of the network and the traffic profiles as well as the abilities of the equipment to change operation state are exploited towards this direction. An indicative example of an energy aware traffic engineering approach is to exploit the dense interconnection of an infrastructure to route traffic over energy efficient paths while putting unused network equipment to sleep. Furthermore, resource allocation strategies [90] and the algorithmic efficiency of the developed protocols [99] directly affect energy consumption and provide the potential for further energy saving.

3.6 Standardization bodies

In this section, a short description is provided for some of the existing standardization efforts on energy efficient networks. Focus is mainly given on fixed access and core networks.

The European Union already published a number of Codes of Conduct covering different categories of equipment, including broadband equipment, data centres, power supplies, UPS. The Code of Conduct on Energy Consumption of Broadband Equipment has been defined by the EU, which sets targets in reducing energy consumption in the access network [108]. This Code of Conduct sets out the basic principles to be followed by all parties involved in broadband equipment, operating in the European Community, in respect of energy efficient equipment. The EU Code of Conduct on Data Centre Energy Efficiency has been also specified in response to increasing energy consumption in data centres and the need to reduce the related environmental, economic and energy supply security impacts [109]. The aim is to inform and stimulate data centre operators and owners to reduce energy consumption in a cost-effective manner by improving understanding of energy demand and recommending best practices.

IEEE has also ratified the Energy Efficient Ethernet (EEE) standard in October 2010, also known as IEEE 802.3az, which is a set of enhancements to the twisted-pair and backplane Ethernet networking standards that will allow for more than 50% less power consumption during periods of low data activity, while retaining full compatibility with existing equipment. The standard is finalized; however its full impact won't be realized for a few years, since new hardware is required to benefit from EEE.

ENERGY STAR is a joint program of the U.S. Environmental Protection Agency and the U.S. Department of Energy that has defined the ENERGY STAR Product Specifications. One of the main goals of the ENERGY STAR program is to develop performance-based specifications that determine the most efficient products in a particular category (i.e. datacentre storage and small network equipment). ENERGY STAR identifies products where large gains in energy efficiency and pollution reduction can be cost-effectively realized and can play an influential role to expand the market for these products.

IETF has recently established the Energy Management (EMAN) Working Group. The EMAN WG will investigate existing standards such as those from the IEC, ANSI, DMTF and others, and reuse existing work as much as possible. The EMAN WG will work on the management of energy-aware devices, covered by the following items:

- Specifying energy management properties that will allow networks and devices to become energy aware.
- Describing extensions to the current management framework, required for energy management. Focus will be put on energy management for IP-based network equipment and the issues of discovery and identification of remote devices.
- Developing a MIB module for monitoring energy-aware networks and devices.
- Defining managed objects for monitoring of power states and energy consumption/production.
- Defining managed objects for battery monitoring, capable of reporting detailed battery state properties and statistics.
- Developing an applicability statement, describing the variety of applications that can use the energy framework and associated MIB modules.

Different interesting issues are under consideration by the Environmental Engineering Technical Body in ETSI: The ETSI EN 300 019 series addresses different topics dealing with environmental conditions and environmental tests. The ETSI EN 300 119-3 is dealing with engineering requirements for cabinets; the ETSI EN 132-3 is exploring novel techniques for power supply interfaces at the input to telecommunications equipment. Other interesting activities in ETSI regard the following topics: ETSI TR 102 489 is related to thermal management of cabinets and discusses how to increase the efficiency of the cooling system in data centres and telecommunication centres; ETSI TR 102 530 covers various methods of increasing the efficiency of Telco systems by controlling/reducing the energy consumption in the TLC network equipment and its related infrastructure. ETSI TR 102 532 is devoted to the feasible usage of alternative energy sources in telecommunication installations; ETSI TR 102 533 defines the power consumption limits, the methodology and the test conditions to measure the power consumption of broadband fixed telecommunication networks equipment. ETSI TR 102 614 is discussing aspects related to reverse powering in fixed access networks. In addition to the previous activities, ETSI ATTM also launched at the end of 2008 a Special Task Force on Efficient Broadband [110], with the objective of producing a number of specifications for the energy efficiency of all segments of the networks (e.g., access networks' equipment, data centres, home gateways and network terminations, etc.).

The Home Gateway Initiative (HGI) launched an internal task force called ”Energy Saving” [111] with the objective of setting up requirements and specifications for energy efficiency in the home gateways, starting from a reference architecture depicted in the HGI home gateway residential Profile v1.0. The analysis could also be extended to the Network Termination (NT) used in case of the ”2box” solution (NT + service router, possible scenario mainly for FTTH) and to other home network infrastructure devices, and it is well linked with the EU Code of Conduct discussions.

Within ITU, ITU-T Study Group 15 (Optical transport networks and access network infrastructures) is the home of the digital subscriber line (DSL) standards that provide the broadband Internet connections for over 600 million households around the world. SG15 standards (ITU-T Recommendations) relating to passive optical networks (PONs) are a crucial step towards all-optical networks, while defining how significant savings may be achieved. Furthermore, ADSL2 and ADSL2+ (ITU-T Recommendations G.992.3 [112] and G.992.5 [113]) already support multiple link rates and power states [114].

In addition to these initiatives, ITU-T created in September 2008 a new Focus Group, namely, FG ICT & Climate Change, with the objective of analysing the positive impacts of ICT on other industry sectors and evaluating actions for optimising the power consumption in the ICT world, as well. Four deliverables have been finalized by March 2009: "Definitions", "Gap Analysis", "Methodologies" and "Direct and indirect impact of ITU Standards". The Focus Group has recognized a number of potential gaps in standards which were acknowledged in the various contributions and recapped in the four Deliverables. Therefore, it has identified a number of activities, which – if followed through as a team effort across the ITU and strongly supported and coordinated by ITU management - could lead to greenhouse gas emission savings in line with existing or emerging targets and timescales.

Finally, ATIS NIPP Specifications on Energy Efficiency define transport products and systems as well as a methodology to calculate the Telecommunication Energy Efficiency Ratio (TEER) of a transport system or network configuration.

3.7 Summary

Section 3 has provided an overview of the existing Internet and the foreseeable evolutions in terms of technologies and traffic characteristics. Based on existing reports, it has been shown that Internet topology is changing from a hierarchical structure to a flatter and more meshed one. This change is followed by the increase in video traffic and the emergence of content distribution networks, as well as by the rise of regional transit ISPs. Traffic volume exchange in the Internet is rapidly increasing, dominated by the Internet video traffic. Predictions report that IP traffic will grow at a compound annual rate of 32 per cent from 2010 to 2015, while the sum of all forms of video will continue to be approximately 90 per cent of global consumer traffic by 2015.

The increase in the traffic volumes is followed by a corresponding increase in the energy consumption of Telcos and ISPs. Due to this trend, 'Going green' is now a necessity for telecom operators in order to reduce energy expenses, conform to restrictions on carbon footprints and emphasize on their corporate responsibility. This necessity is also confirmed by the Global e-Sustainability Initiative reports that estimate the carbon footprint of networks and related infrastructures at about 320 Mtons of CO_2 emissions in 2020. As far as wired networks are concerned, both telcos' devices and broadband access equipment will cause ever growing and non-negligible contributions, equal to 22% and to 15%, respectively. Furthermore, capacities of high-end routers grow in exponentially, by a factor of 2.5 every 18 months, while at the same time silicon technologies improve their energy efficiency with a factor of 1.65 every 18 months.

Taking into account the increasing demand for energy efficiency, several next generation technologies that can be proved energy efficient have been described. Optical backbone networks can exploit the enormous bandwidth that may be provided by fibre optical technologies and the flexibility for establishment of end-to-end optical circuits among various nodes in the network with reduced energy demands. Content-centric networking is an alternative and innovative approach to the architecture of computer networks that can reduce transit traffic in the backbone network links and finally lead to improved energy efficiency. Autonomic networking is able to offer a basic platform to achieve greener Internet services, since it looks for better resource management and utilization according to the status of the system, while Cloud computing can contribute to the reduction of energy consumption for processing and storing data, since part of computing and storage resources may be highly utilized while the remaining part can be in sleeping status. Furthermore, next generation technologies in the access part have been also presented with a categorization according to the supported access bit rate and the corresponding power consumption.

Then, the energy aware design space has been introduced for the device data plane, as well as the control plane. The device data plane refers to re-engineering, dynamic adaptation and sleeping/standby approaches, while the control plane includes energy aware traffic engineering mechanisms, virtualization techniques and efficient resource allocations approaches. Finally, the most important standardization bodies and groups in the field of green networking and green ICT have been shortly presented.

4 Energy efficiency requirements for next generation networks and devices

In this section, information is provided regarding energy consumption in today's networks and devices. Requirements for energy efficiency support are described, as well as target energy saving for each element of next generation networks devices upon the application of energy-aware technologies. Furthermore, three reference network scenarios are presented in detail, focusing on energy consumption on current networks. Energy consumption is measured per network part and the effects of traffic volumes variation (diurnal patterns, traffic volumes per network segment, average loads) on real time energy consumption are examined. The reference scenarios refer to a telecom operator in Italy and two ISP providers in Greece and Poland. In each case, future energy aware technologies that may be used in order to minimize energy consumption are identified.

4.1 Energy consumption in today's network devices

In order to face the energy efficiency issue in today's and tomorrow's wire-line networks, we have to firstly understand and accurately characterize the real sources of power wastes. As outlined in this section, network devices working in the different network portions play a central role, since the overall energy consumption in networks arises from their operational power requirements and their density.

High-end IP routers based on multi-chassis platforms can be certainly considered as the network node typology with the highest complexity level. Traffic processing engines have generally to support complex forwarding and lookup functionalities, and internal HW elements are generally dimensioned for processing enormous traffic volumes. So, green technologies will have to mainly address the energy-efficiency in packet processing engines in order to effectively reduce the carbon footprint of such kind of devices.

On the other hand, Digital Subscriber Line Access Multiplexers (DSLAMs) generally include less and simpler packet processing functionalities with respect to routers, but also present a much larger number of network interfaces. The current energy requirements of today's DSLAMs mainly spring from link interfaces. Therefore, future green technologies and solutions for access networks will have especially to focus on energy-efficiency at the link/network interface level.

Evaluating the current status of power consumption for the access devices working at the different network levels can help to provide a true characterization of the energy needs for the current architectures and then estimate the future needs taking into account the technology evolutions and the implementation of mechanisms for better efficiency. To this purpose, the following subsection will try to characterize the energy-efficiency and requirements of devices working at the access, home, metro/transport and core networks, respectively. Data are related to products already on field or prototypes tested in the Labs and are referred in general to equipment not designed with specific attention to energy efficiency issues, but widely deployed. In some cases, products manufactured following specific energy-related requirements are considered.

4.1.1 Access networks

The access network equipment taken into account in this chapter deals with the Fixed Broadband network. In particular, such equipment includes:

- MSAN (Multi Service Access Node), which can be installed in central offices as well as in FTTCab architectures. Such devices have typically 14-20 slots which can host different types of line cards (ADSL/ADSL2+, VDSL2, POTS, ISDN, SHDSL, GPON, GbE), resulting to a total number of 600-1200 BB ports.
- OLT (Optical Line Termination), which are typically installed in COs (Central Offices) for FTTH (Fibre To The Home) deployments. Such equipment can support both GPON and PtP architectures.

The power consumption for 11 different configurations of the above equipment is reported in the Figure 29 - Figure 39. Such power consumption is defined as:

$$P_{BBline} = P_{BBeq} / N_{subscrib\text{-}lines} \quad \text{where:}$$

- P_{BBeq} is the power consumption (in W) of the fully equipped Broadband equipment (DSLAM, OLT), measured at the electric power input interface, placed at the premises of the operator or the equipment supplier, which connects multiple broadband subscribers to a backbone.
- P_{BBline} is the power consumption per line in W of the Broadband equipment.
- $N_{suscrib\text{-}lines}$ is the maximum number of subscriber lines served by the broadband equipment (DSLAM, OLT) under test.

The power consumption measured in all the 11 configurations refers to full power mode. Moreover, the following items have been considered part of the DSLAM/OLT and therefore their power consumption has been taken into account to get the total power consumption (P_{BBeq}) of the equipment:

- Network Termination board, providing one or more links to the Core or Backhaul Network.
- Line Termination board, providing a number of xDSL/POTS/GPON ports connected to the end-user through the local loop.
- Splitter (Low Pass Filter) function.
- Backplane (or other) to interconnect the different blocks of the DSLAM/OLT.
- Inside Rack Cooling system (e.g. fans drawer inside cabinet based DSLAM/OLT systems).
- Normal operational power supply unit.

The following items have not been considered part of the equipment and therefore their power consumption has not been added to the power consumption values:

- Rectifier (AC/DC).
- Room or outdoor Cabinet Ventilation and Air Conditioning Unit (VAC Unit).
- Auxiliary power unit.

- Battery.
- Additional External signal processing (Dynamic Spectrum Management (DSM) and Multiple-Input Multiple Output (MIMO) techniques if not implemented as part of the Line Termination board).

The power consumption measurements for DSLAM equipment described in the present document apply at Interface "A" as shown in Figure 26.

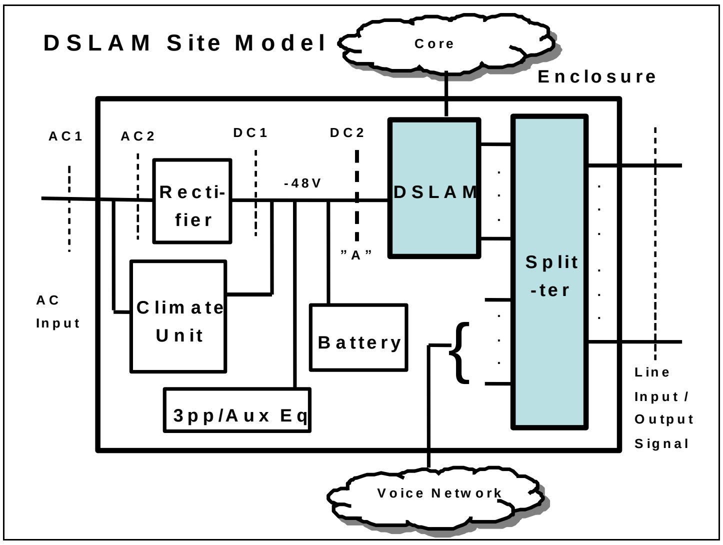

Figure 26: DSLAM node site reference model.

The basic test setup for DSLAM which has been used during the power measurements tests is reported in Figure 27**.** Both the network side (optionally through an Ethernet switch) and the end-user side (direct or also through an Ethernet switch) are connected to an Ethernet Traffic Simulator/Analyser.

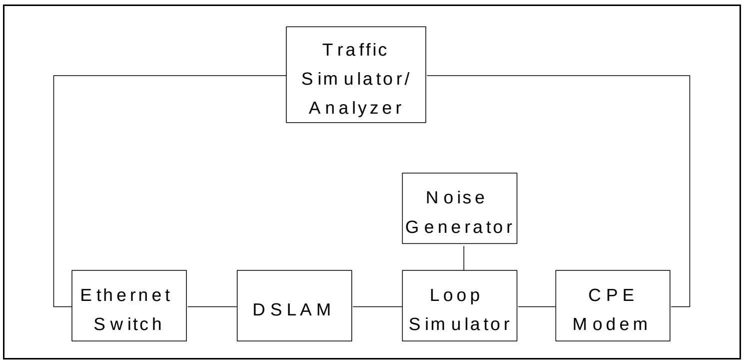

Figure 27: Test Setup for power measurement for DSLAM.

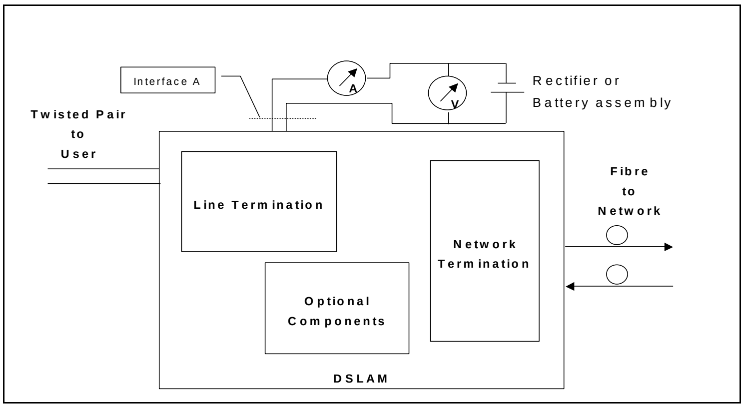

Figure 28: Power Consumption at System level.

In Figure 28, the actual DSLAM power measurement method is shown. The DSLAM comprises of the line termination boards, the Network termination boards and some other components like the cooling system. The Network termination board has fibre connections to the traffic simulator/analyser (as shown in Figure 27) and the line termination boards have twisted pairs connected to loop/line simulators. The power of the system is measured at the "A" interface of the DSLAM using both a current and voltage meter. The system can be powered either through a battery assembly or rectifier set at the nominal voltage (for the actual test, the second has been used).

Even if the most precise and correct way to measure P_{BBeq} is to have the DSLAM fully equipped and all its ports connected to CPE through loop or line simulators, it must be recognised that this is too complicated and time consuming. Therefore, for the 11 considered configurations, an alternative technique has been used that reduces the number of CPE required and requires extrapolation to give the correct per line result. In particular, in a first step the power consumption P_{empty} of the equipment, with just common parts (Network Interface Cards) and without any Line Card, has been measured. Then, in a second step, one Line Card has been added to the system with all lines connected using the Show Time Freeze functionality. The power consumption of the DSLAM with the added Line Card ($P_{1\ line\ card}$) has then been measured once again and the difference ($P_{1\ line\ card}$ - P_{empty}) has been considered as the power consumption ($P_{line\ card}$) of a fully equipped Line Card. Finally, the total power consumption has been calculated according to the following formula:

$$P_{BBline} = P_{empty} + n * P_{line\ card}$$

where n is the maximum number of Line Cards per DSLAM. A similar process has been followed to evaluate the unit power consumption of a fully equipped GPON OLT. In particular, also in this case a first step aimed at measuring the P_{empty} of the equipment has been performed. Then, the $P_{line\ card}$ has been measured connecting one ONT for each of the 4 or 8 GPON ports of a single Line Card. The extrapolated total and unit power consumption for each of the 11 different configurations are reported in the following figures.

All the figures have a similar structure: there is a red line, continuously growing in steps, which represents the trend of the total extrapolated power consumption of the equipment (in terms of W - the associated values are reported on the left vertical axis of the chart) with respect to the progressively increasing number of its active ports (ADSL2+, VDSL2, POTS or GPON). It is important to highlight that such red line does not start from 0 W, even if no line towards the final customer is active. This is because such equipment has the so called “common parts” (usually a couple of network line cards for redundancy), which contain the uplink ports towards the metro/core network, as well as the switching matrix, and are always active (and therefore consuming electrical energy, typically 50 W-80 W for each card). Then, the presence of the "steps" is strictly related to the modularity of the customer line cards (for example, in case of ADSL2+, the typical configuration foresees 48-72 ports for each card). More precisely, the total extrapolated power consumption grows linearly as long as each customer line card is progressively filled: once that such line card is completely full, it is assumed to add another line card, which itself has a certain fixed energy consumption even if all its ports are still not active. In conclusion, the steps are due to the "common parts" of the single customer line cards progressively supposed to be added.

The charts then report an orange line, similar to the red one, which instead represents the trend of the total extrapolated power consumption of the equipment (in terms of W) with respect to the progressively increasing number of its inactive (non-active) ports. In practice, the steps are due to the insertion of a further customer line card within the equipment, but such steps are then flat and not linearly increasing, just because it is assumed that no port is active and therefore consuming energy for carrying data traffic. The aim of the orange line is then to show the final maximum gap in terms of power consumption between a piece of equipment with all its ports actively carrying data traffic and the same equipment with all the customer line cards inserted but all the ports not activated. As an example, as far as the MSAN in Figure 29 is concerned, the gap between the two configurations described is roughly 300 W. For a TLC Operator, such gap represents important information because in a real deployment it is of course not realistic to assume that newly installed BB equipment will immediately have all its ports connected to active customers: therefore, its power consumption will stay between the red line and the orange one. The information brought by such lines will then support a better estimation of the trend of the BB network's total power consumption with the increasing number of service subscribers.

Finally, each following chart reports a blue line, which instead progressively decreases with the growing number of active ports. This is because such line represents the unit power consumption (in terms of W/port) of the equipment; the associated values are reported on the right vertical axis of the chart. For example, concerning Figure 29, initially the unit values are quite high (around 8W/port): this is of course because the power consumption of the already mentioned common parts is divided by a low number of active ports. Then, as long as more customer line card are added, more lines are activated and therefore the unit value progressively decreases (the "steps" are always due to the small common parts of each single customer line card). With the fully equipped configuration, the unit power consumption of the equipment in Figure 29 drops to 1.21 W/port.

Where available, the figures also indicate in the dashed box the target unit power consumption reported in the European Commission BB Code of Conduct (CoC, see paragraph 3.6). For example, concerning Figure 29 the MSAN ADSL2+ reference value is 1.35 W/port. This brings us to the conclusion that some current equipment has already reached a good level of unit power consumption (as in many cases it is lower than the associated CoC target). Anyway, further improvements are needed provided that the reported targets do not take into account innovative solutions (e.g. low power modes) aimed at reaching better energy efficiency performances.

In Figure 29, the equipment can host a maximum number of ADSL2+ ports of 384. The figure reports the results of a mix of power consumption measurements and extrapolations. In particular, the power consumption of the common parts and of one single line card (loop back-to-back) has been measured. Then, through the extrapolations, the total power consumption (as well as the unit power consumption) of the fully equipped configuration has been estimated.

In Figure 30, the equipment can host a maximum number of ADSL2+ ports of 1280 while in Figure 31, the equipment can host a maximum number of ADSL2+ ports of 1024. The figures report the results of a mix of power consumption measurements and extrapolations, as described for the first configuration tested.

In Figure 32, Figure 33 and Figure 34, the equipment can host a maximum number of POTS ports of 384, 1280 and 1024 accordingly. The figures report the results of a mix of power consumption measurements and extrapolations, as described for the first configuration tested.

In Figure 35, the equipment can host a maximum number of SHDSL+ ports of 192 while in Figure 36 the equipment can host a maximum number of SHDSL+ ports of 1280. The figures report the results of a mix of power consumption measurements and extrapolations as in the previous cases.

In Figure 37, Figure 38 and Figure 39, the equipment can host a maximum number of 128 GPON ports. The figures report the results of a mix of power consumption measurements and extrapolations as in the previous cases.

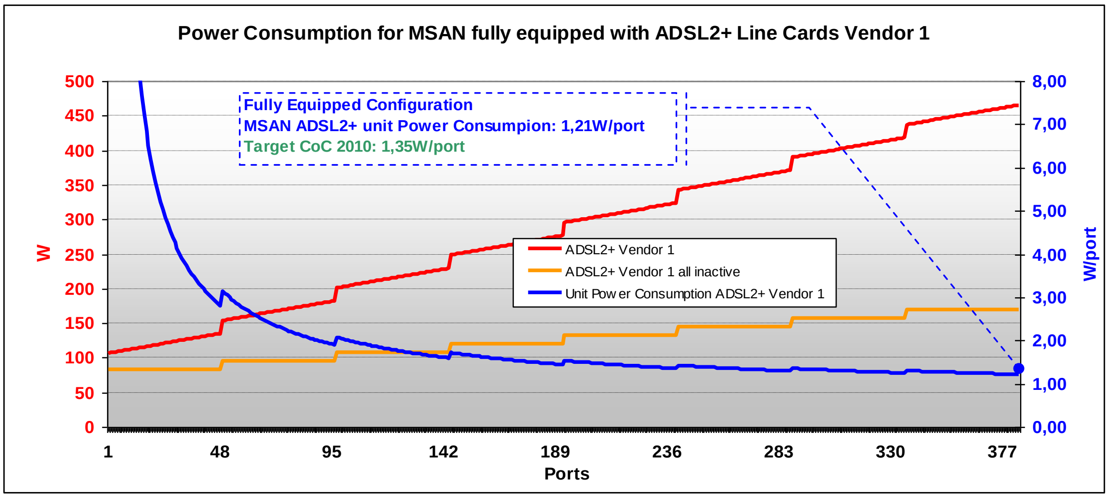

Figure 29: Multi-service access node (MSAN) fully equipped with ADSL2+ line cards (Vendor 1).

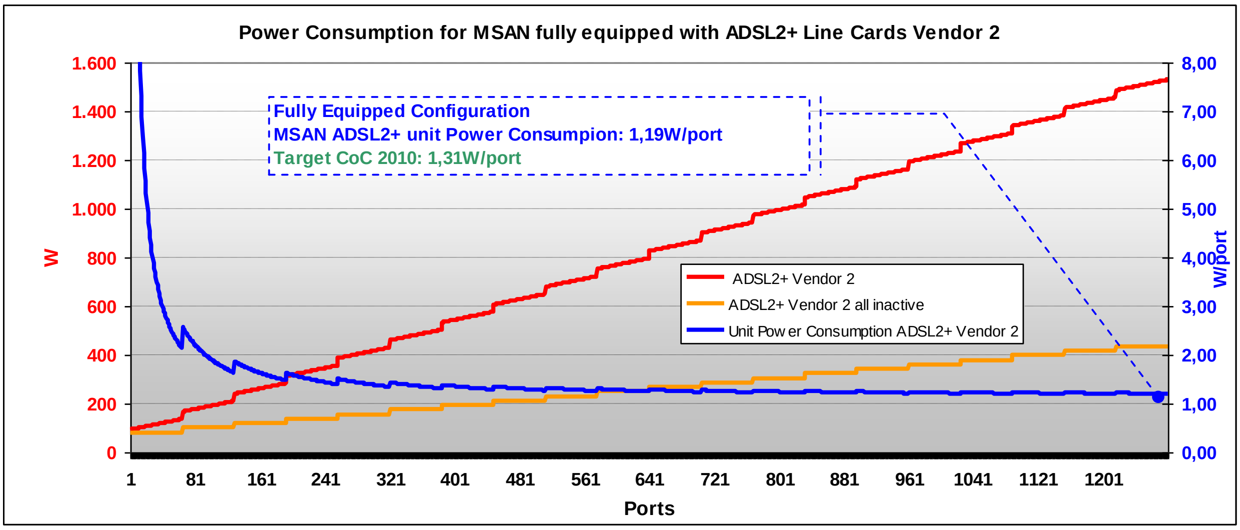

Figure 30: Multi-service access node (MSAN) fully equipped with ADSL2+ line cards (Vendor 2).

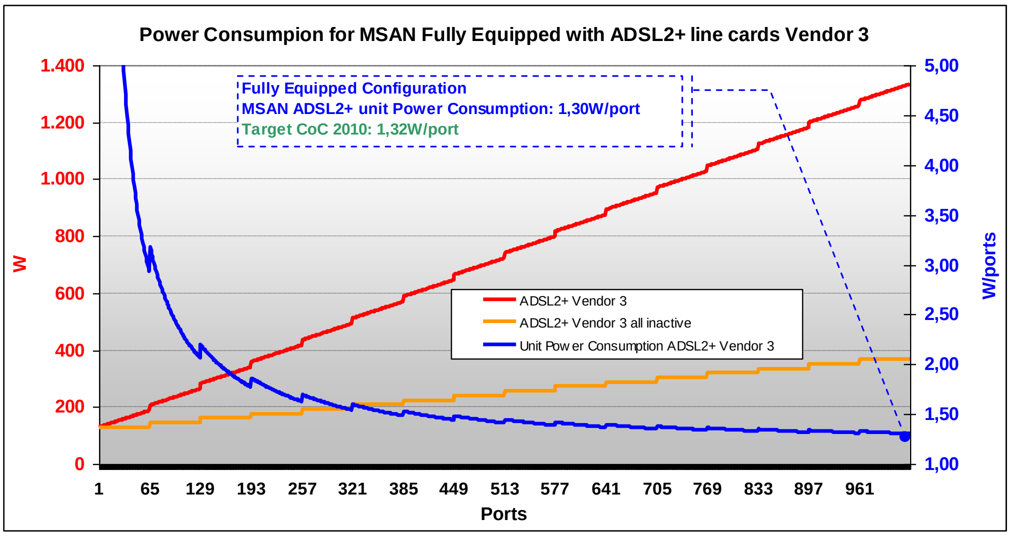

Figure 31: Multi-service access node (MSAN) fully equipped with ADSL2+ line cards (Vendor 3).

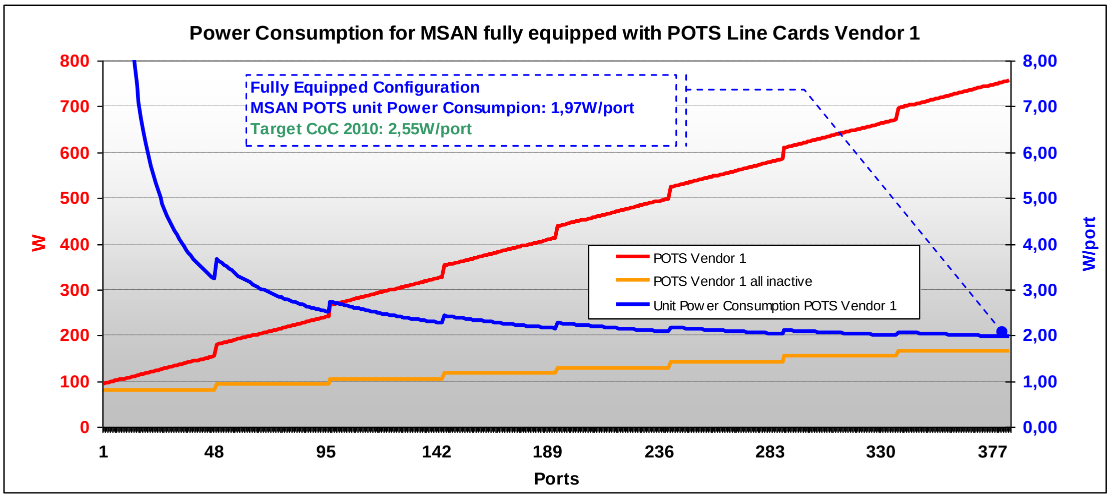

Figure 32: Multi-service access node (MSAN) fully equipped with POTS line cards (Vendor 1).

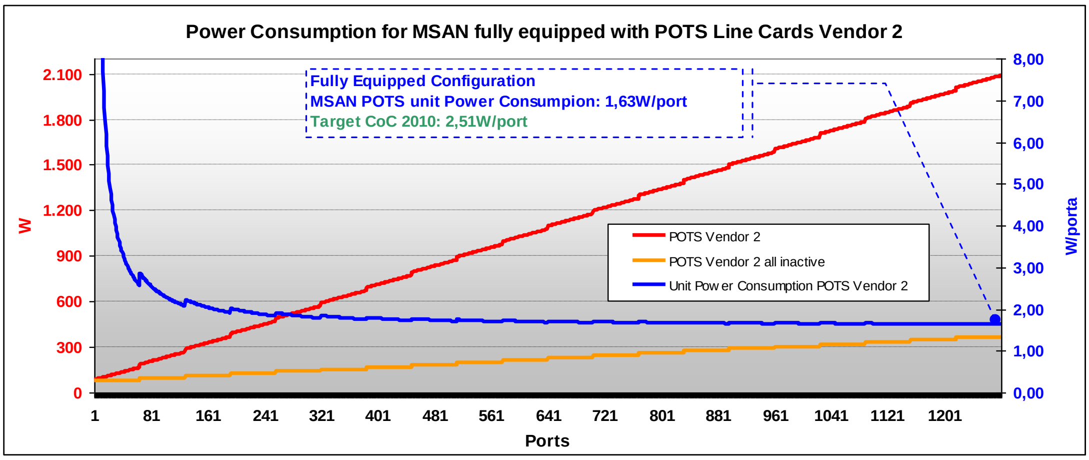

Figure 33: Multi-service access node (MSAN) fully equipped with POTS line cards (Vendor 2).

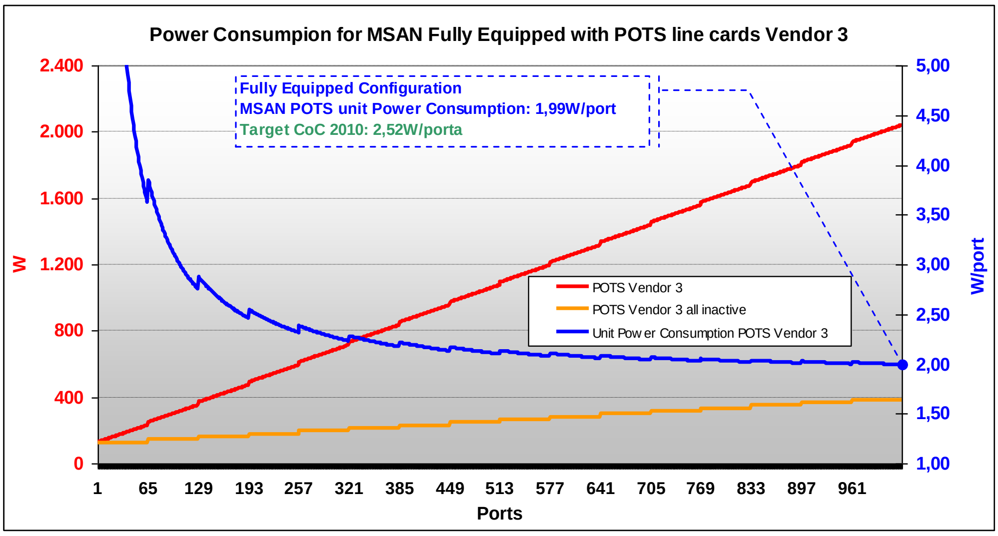

Figure 34: Multi-service access node (MSAN) fully equipped with POTS line cards (Vendor 3).

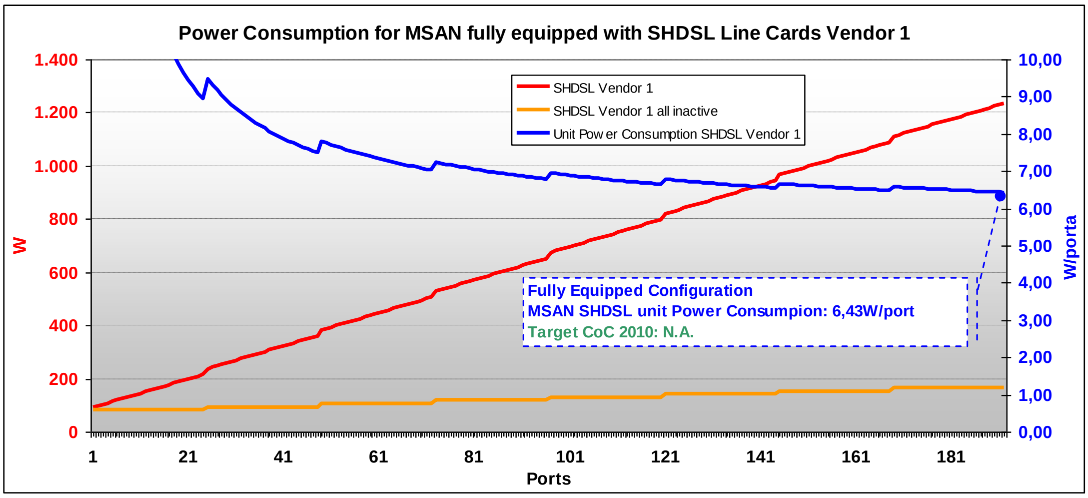

Figure 35: Multi-service access node (MSAN) fully equipped with SHDSL line cards (Vendor 1).

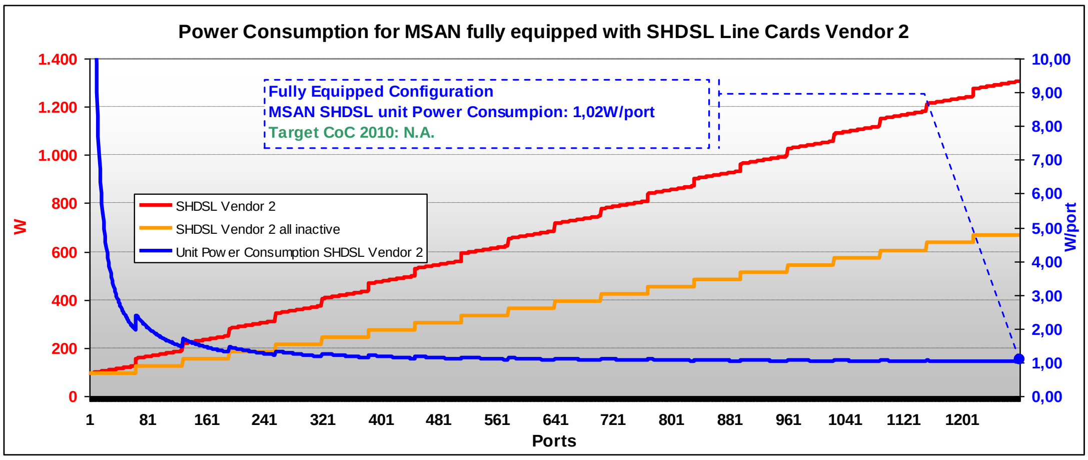

Figure 36:Multi-service access node (MSAN) fully equipped with SHDSL line cards (Vendor 2).

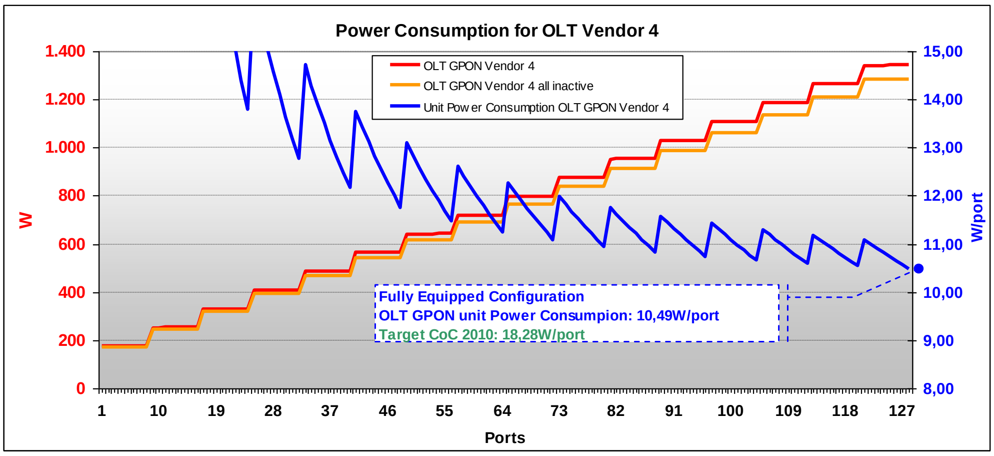

Figure 37: Optical Line Termination fully equipped with GPON line cards (Vendor 4).

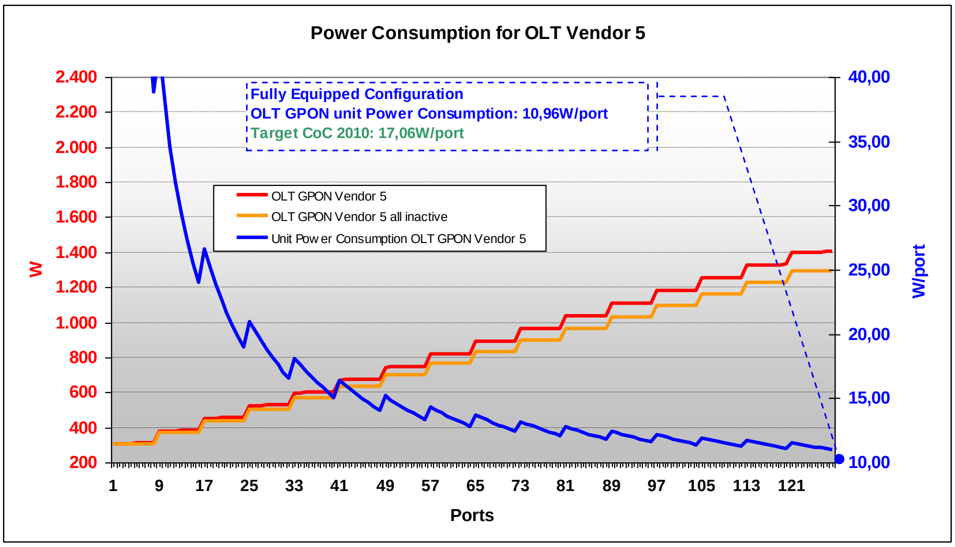

Figure 38: Optical Line Termination fully equipped with GPON line cards (Vendor 5).

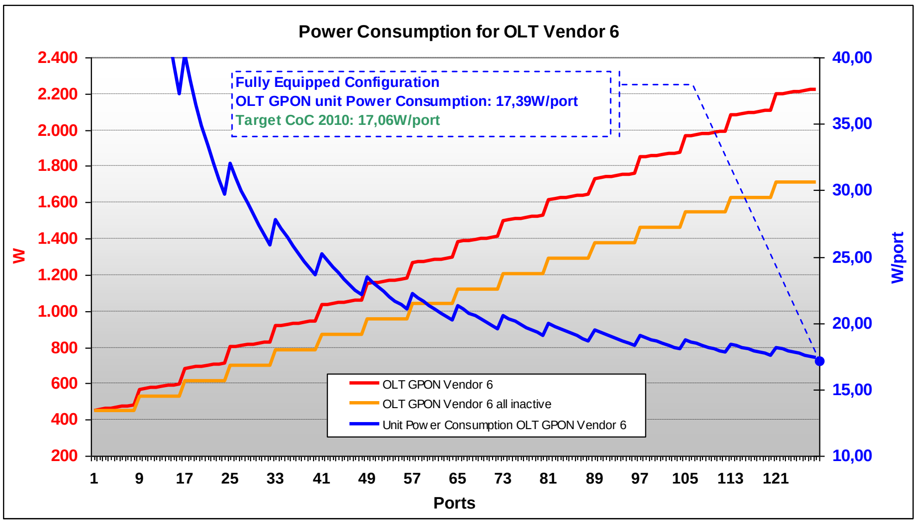

Figure 39: Optical Line Termination fully equipped with GPON line cards (Vendor 6).

Table 3 depicts the main information excluded from the power consumption measurements reported in the previous figures. In particular, it contains information regarding the type of the analysed equipment, the related vendor, the specific technology, the nominal peak rate per port of the specific technology, the power consumption per port in a fully equipped configuration, an indicator of the ratio between the nominal peak rate per port and the associated measured power consumption (expressed in terms of Mbit/s per Watt) and the unit power consumption target reported within the European Commission Code of Conduct on Broadband equipment 2010 [108].

Table 3: Comparison of the energy consumption measurements.

Equipment	Vendor	Technology	Peak rate per port	Power consumption per port	No of users	Power consumption per user	Peak rate x users/power [2] [Mbit/s/W]	CoC target BB 2010
MSAN	1	ADSL2+	20 Mbit/s [3]	1,21W	1	1,21W	16,53	1,35W
MSAN	2	ADSL2+	20 Mbit/s [3]	1,19W	1	1,19W	16,81	1,31W
MSAN	3	ADSL2+	20 Mbit/s [3]	1,30W	1	1,30W	15,38	1,32W
MSAN	1	POTS	64 kbit/s	1,97W	1	1,97W	0,032	2,55W
MSAN	2	POTS	64 kbit/s	1,63W	1	1,63W	0,039	2,51W
MSAN	3	POTS	64 kbit/s	1,99W	1	1,99W	0,032	2,52W
MSAN	1	SHDSL	2Mbit/s [3]	6,43W	1	6,43W	0,311	N.A.
MSAN	2	SHDSL	2Mbit/s [3]	1,02W	1	1,02W	1,96	N.A.
OLT	4	GPON	2,5Gbit/s	10,49W	32 [4]	0,33W	7626	18,28W
OLT	5	GPON	2,5Gbit/s	10,96W	32 [4]	0,34W	7299	17,06W
OLT	6	GPON	2,5Gbit/s	17,39W	32 [4]	0,54W	4600	17,06W
OLT	- [5]	Pt-Pt	1Gbit/s		1		416	2,4W

Based on the values reported in Table 3, the following conclusions are extracted:

- The vast majority of the equipment measured has shown a good compliance with the power consumption targets set in the EC Broadband Code of Conduct. This could mean either that the current market products have reached a good level of energy efficiency, or that the power consumption targets for 2010 were not so challenging. The truth stands probably in

[2] The indicator "Peak rate x users/power [Mbit/s/W]" has been chosen as the most capable to really represent the real efficiency of the networks. Its use of the number of users is based on the assumption that, within the wireline transmission systems, the physical multiplexing present in point to multipoint transmission systems (e.g. in passive optical network systems) is equivalent to the traffic multiplexing present on all point to point systems where the aggregated speed at the uplink interfaces is always much smaller (1/10-1/50) than the aggregated speed of all the user interfaces. In this way, such an indicator helps in clearly highlighting that the PON is definitely more efficient than the Point to Point (PtP) architecture, since it can provide service to several customers at the same time and without any relevant bandwidth difference encountered by the customer themselves. The peak rate and the power within the indicator refer to the single port.

[3] The line speeds for the DSL systems (SHDSL, ADSL and VDSL) vary heavily depending on the line length and can be typically far lower than the peak declared value. SHDSL systems can show a limited reduction vs line length, while A/VDSL mean speed lies at about 1/3 to 1/5 of the peak speed.

[4] Although GPON equipment are point to multipoint systems capable to serve up to 128 users per each interface, a mean amount of 32 users per interface has been used to better represent real life scenarios.

[5] Point to point optical interface consumption has been derived from the EC Broadband Code of Conduct.

the middle, even if it must be highlighted that the same analysis performed few years ago were absolutely not showing such a wide percentage of compliance. This means that certainly a good progress has been performed in the field of the energy efficiency of broadband equipment. Moreover, the new targets for the period 2011 – 2013 reported in the CoC BB V4 will push towards further improvements.

- There is a huge difference between the energy efficiency of broadband equipment related to the same time period but to different technologies. To this end, if we take into account the nominal peak rate per port divided by the related power consumption, we find that the ADSL2+ technology is ten times more efficient than the SHDSL one. Moreover, at the same time the GPON technology is more than one hundred times more efficient than the ADSL2+. This of course confirms the absolutely better performances of the optical fiber with respect to the traditional copper pair. The table does not report any measurement of the VDSL2 technology that can reach speeds of 50-100Mbit/s with a related CoC target of less than 2W/port. Such technology is therefore two to three times more energy efficient than ADSL2+ and hence slightly closer to the optical fiber performances; nevertheless, it must be highlighted that also the GPON technology is evolving, and the already commercially available XG-PON1 can guarantee four times the speed at less than double the power. Moreover, no need to highlight the remarkable flexibility offered by the GPON architecture, which currently allows up to 128 customers to be connected to the same optical GPON port. Such splitting ratio will evolve to up to 1:512 with the future NG-PON2 solution. Finally, according to the EC BB CoC, optical point to point 1Gbit/sec systems show energy consumption per line in the same range as V/ADSL systems and a Peak rate/power ratio much higher than the DSL systems, but ten times lower than the GPON systems.

4.1.2 Network terminations

Home Gateways (HGs) are the devices that provide network termination functionalities and, thus, access to the Internet and its services (Figure 40). Voice-over-IP, standard telephony POTS (Plain Old Telephony System), DECT, wireless and Ethernet connectivity, and USB additional peripherals, like printers, mass storage devices are required in a HG, that provide multi-standard and always available services [101]. This implies the HG should be powered on during the entire day, acting like an always on communication device. However, as the data reported in Figure 41, for the largest part of the day the HG is in standby mode confirming that access to the HG features is not continuous. The off period is usually during the night. The 0 W period consists when the power supply is also powered off. Furthermore, Figure 42 shows that the energy consumption caused by the power supply is actually the largest part of the overall energy consumption in a home gateway.

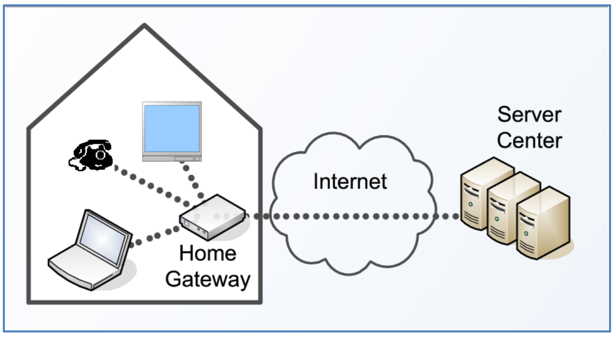

Figure 40: Home Gateways as the central home communication device.

Taking into account the above-mentioned HGs usage characteristics, in this subsection, results of measurements in Telecom Italia Labs related to network termination in the form of a home gateway or optical NT are provided. Network terminations are considered here as part of the access network, but they are in fact on the boundaries between that and the home environment. In fact, the equipment considered can implement not only functions related to the network termination as such, but also other features allowing the service usage within the home network, from connectivity to service layers.

There are two basic architectural configurations to be considered. The first one is including all the connectivity functionalities (towards the WAN and the LAN side) and service support in one box only, called as "home gateway" in Figure 43. A second possible configuration implements a "two boxes" model and is typically used in next generation networks with fibre connection on the WAN side, as it is shown in Figure 44. Both figures describe in more details the scenarios that the measurements are referred to.

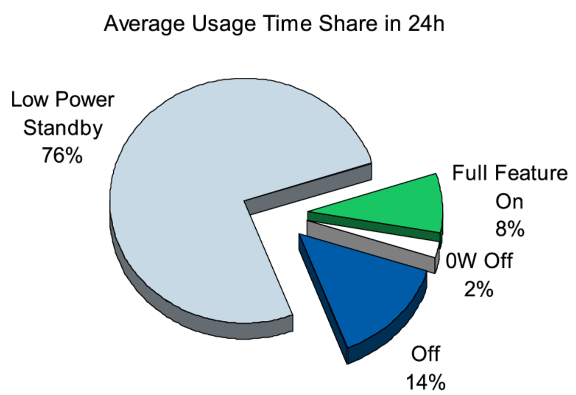

Figure 41: Average usage time share in 24 h for HGs.

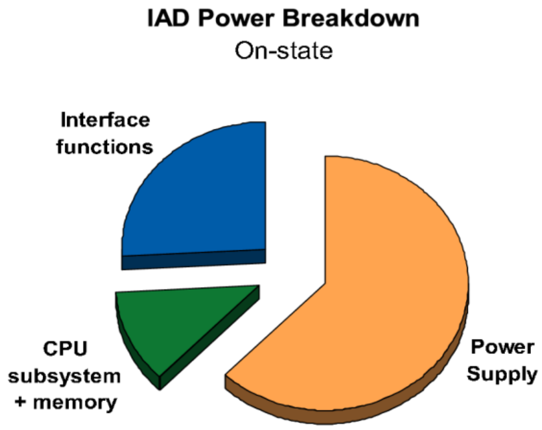

Figure 42: HG power consumption share for non-optimized power system, CPU subsystem and interface functions.

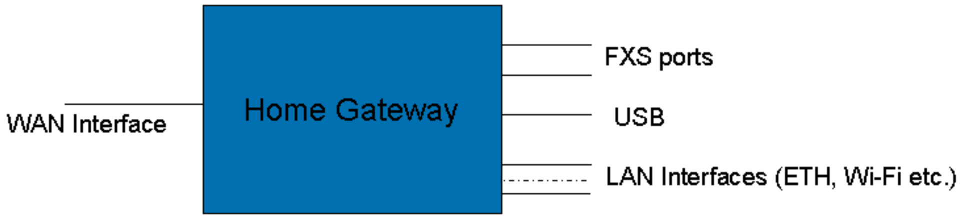

Figure 43: Case 1: xDSL termination with a single box (Home gateway).

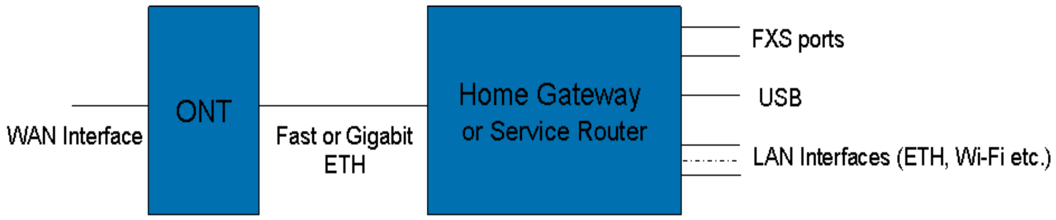

Figure 44: Case 2: GPON termination with a double box (Optical NT + Home gateway).

For these categories of equipment, defining the operational states subject to measurements is also important, as they can enter different modes of operations during the day depending on the number and type of service supported. Different power levels are therefore associated to every single state. The reference for the definition of states considered here is the European Code of Conduct for Broadband Equipment, covering this category of devices. The Version 3 of the CoC used as reference for these measurements, is defining two basic states, the "low power" and the "ON/full power", as outlined in Table 4. It should be noted that even if the Code of Conduct mentioned in this list of states is Version 3, the new Version 4 has been produced and published in February 2011, but the impact on the operational states definition is really low, while there is a significant update of targets that ECONET will take into account in terms of achievement of challenging energy efficiency objectives.

Table 4: Definition of energy-aware states according to the version 3 of the CoC.

Low Power States	
Port/component	*Low-Power State*
Central Functions (processor and memory: routing, firewall, OAM (e.g. TR-069), user interface)	Not processing user traffic
WAN interface	Idle (link established, but no user traffic transmission)
LAN Ethernet ports	Ports not connected (or no Ethernet link) but with Ethernet link detection active
Wi-Fi	Beacon on, but no user traffic transmitted
Alternative LAN technologies (HPNA, MoCA, Powerline, POF, etc.)	MoCA, Powerline, HPNA or POF capability is activated, but no user traffic transmitted
FXS	1 phone connected, on hook. Off hook detection active. If there are multiple FXS ports, only one is connected
FXO	No active call, incoming call detection enabled
DECT interface	No active call, incoming call detection enabled
DECT charging station for DECT handset	DECT handset on cradle, in trickle charge
USB	No devices connected, detection of USB devices active

Full Power States	
Port/component	*On State*
Central Functions (processor and memory: routing, firewall, OAM (e.g. TR-069), user interface)	Processing typical level of user traffic
WAN interface	Active (link established at full line rate and passing user traffic)
LAN Ethernet ports	All 4 ports active (link established at max rate and passing user traffic), cable length=5m
Wi-Fi	Beacon on, with user traffic
Alternative LAN technologies (HPNA, MoCA, Powerline, POF, etc.)	MoCA, Powerline, HPNA or POF capability is activated, with user traffic transmitted
FXS	1 phone connected, off hook, 1 active call. If there are multiple ports only 1 is connected.
FXO	1 active call
DECT interface	interface 1 active call
DECT charging station for DECT handset	handset not on cradle, no charging
USB	No USB device connected, detection of USB devices active

The available measurements for the network terminations or home gateways have been further elaborated in order to evaluate the impact on the typical daily duty cycle of these products, defined as follows (as per the contribution, customised for ECONET scenario and context, provided by Telecom Italia to ETNO Energy Task Force in the document "Green Benchmark requirements" published in 2009 - http://www.etno.be/LinkClick.aspx?fileticket=czK6zj6x8yg%3d&tabid=2260)

$$E_{total} = (20 * E_{CoCLOWPOWER} + 4 * E_{CoCON})$$

This parameter gives the possibility of evaluating the average consumption for the single day and then for the single year, for a given product. In the following, the summary of the measurements related to home gateways in the one-box scenario is presented. Commercial products have been tested in their basic configuration: ADSL2+ interface, ETH switch, Wi-Fi g or n, USB. In addition, with reference to the two-boxes scenarios, a number of optical network terminations have been tested, in different configuration as listed in Table 6.

In summary, the two tables constitute the characterization of the network part dealing with the interface between the home environment, where services are actually used, and the access network equipment. In both cases (home gateways and network terminations) a mix of old designs and new energy efficient solutions is present, but the current situation on field (Jan. 2011) is well reflected by the average of the listed values, as the replacement of the old solutions is still on-going. Additionally, the best-in-class products do not yet represent the maximum efficiency that can be achieved, as the implementation of automatic power management mechanisms is still far from a huge deployment.

Table 5: Measurement results for the home gateways in the one-box scenario.

Product	Low Power [W]	Full Power [W]	Yearly Duty Cycle [kWh]
Product A	9.1	11.7	83.51
Product B	5.6	7.7	52.12
Product C	5	6.9	46.57
Product D	8.7	10.7	79.13
Product E	8.9	10.6	80.45
Product F	10.8	11	94.90
Product G	9.8	11.4	88.18
Product H	9.4	11.6	85.56
Product I	6.2	7.5	56.21
Product J	5.5	8.7	52.85
Product K	6.5	9.5	61.32
Product L	5.3	7.5	49.64
Product M	7.4	8.8	66.87
Mean value	**7.55**	**9.51**	**69.02**

Table 6: measurement results for the home gateways in the two-boxes scenario.

Product	Configuration	Low Power [W]	Full Power [W]
Vendor 1	GPON interface, GbE switch with 4 ports	5.4	10.4
Vendor 2	GPON interface, FE switch with 4 ports and 2 FXS ports	5.5	10.1
Vendor 3	GPON interface, GbE switch with 4 ports and 2 FXS ports	11	15,9
Vendor 3 – Simple Configuration	GPON interface and 1 GbE port	8.1	11.1

4.1.3 Switching systems

Switch systems consume much more energy than end nodes. The power consumption of a switch is comprised of the following elements:

- **Number of ports** - Number of ports has a direct contribution to the power consumption of the switch. The power consumption per port is generally fixed and the total consumption has linear relation to the number of ports.
- **Speed / width** - InfiniBand technology supports several lanes (width) of data per port and several speeds per lane. The total bandwidth of each port is (width x speed). The total power per ports increases as the speed and the number of lanes increases but not in a linear way.
- **Chassis size** - Chassis size is derived from the number of ports the switch supports. As the chassis enlarges, the cooling system needs to be more massive and as such it consumes more power. The chassis power supply utilization can also be considered as a power consumer.
- **Active / Passive Cables** - The type of cable has a direct influence on the power consumption of the switch (see section 4.1.5 for more details).

- **Managed / Un-managed** - Switch systems can come with an additional management module that provides chassis and network management capabilities. A management module generally consists of a CPU and memory components that add to the total power consumption of the switch. Some switches have a redundant management node to provide high availability capabilities.

To understand the different elements of power consumption in a switch system we can examine the power calculation of the Mellanox IS5600 648-port InfiniBand Chassis switch (Figure 45) and the Alcatel-Lucent 1850 TSS-160 that is a next-generation, Packet-Optical Transport platform that supports any mix of traffic, from all-circuit to all-packet (Figure 47).

An IS5600 system can support up to 648 external ports; each group of 18 ports has its own node called a leaf. Additionally, it includes 2 management nodes and 8 fan units. The system provides full inter connectivity using spine units. Total max power consumption of a switch is calculated as follows: 1320.23W (base power consumption of the chassis including full power supplies) + (82W * # of spines) + (20W * # of management modules) + (120W * # of fans) + (82W * # of leafs for passive cables) or (121.45W * # of leafs for active cables).

Figure 46 below summarizes the maximum power consumption of different Mellanox IS5XXX InfiniBand switches under full connectivity and full load. As seen in the figure, the major power consumers are the leaf and the spines. Most of the power in these modules is consumed by the network processor and it depends on the temperature, port speed and width. The next big consumer is the use of active cable and then the cooling system.

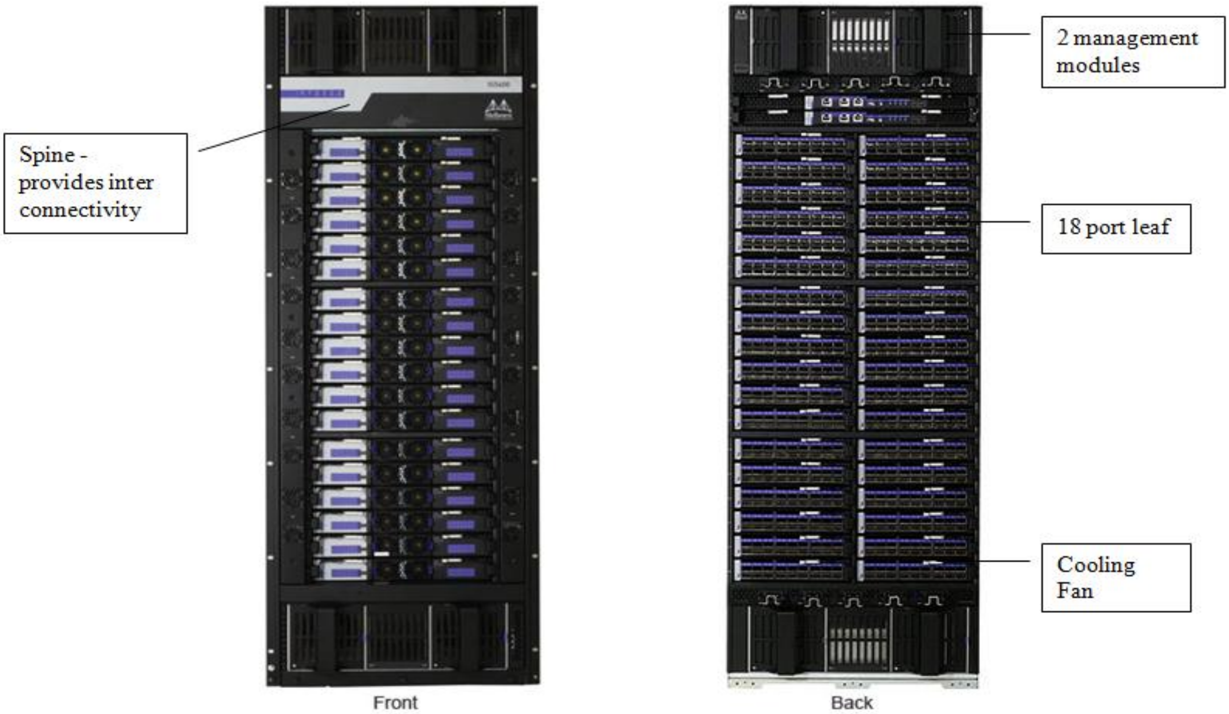

Figure 45: Mellanox IS5600 648-port InfiniBand Chassis Switch.

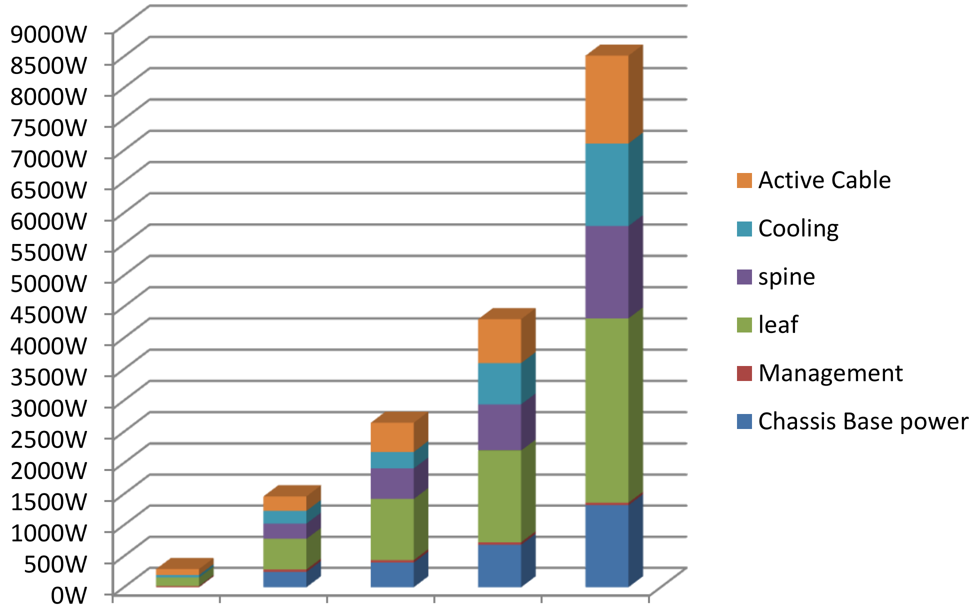

Figure 46: Power Consumption in IS5XXX Family.

Regarding the Alcatel-Lucent 1850 TSS-160, businesses can begin with circuit-based transport and gradually ramp up packet transport by simply changing line cards. The Alcatel- Lucent 1850 TSS-160 supports current traffic requirements while eliminating the scalability issues encountered when traditional multiservice provisioning platforms (MSPPs) are confronted with the high growth of packet-based traffic.

Figure 47: Alcatel-Lucent 1850 Transport Service Switch 160.

The Alcatel-Lucent 1850 TSS-160 offers the flexibility to split increasing traffic demands among any combination of Carrier Ethernet, Transport Multi- Protocol Label Switching (T-MPLS), wavelength division multiplexing (WDM), optical data unit (ODU) and SDH/SONET transport technologies. It offers powerful cross-layer network management and a unified control plane, simplifying operations and reducing the total cost of ownership. Alcatel-Lucent 1850 TSS-160 sub-rack is composed of:

- 8 slots, 20 Gb/s per slot: 16 half slots, 10 Gb/s per half slot
- Two 160 Gb/s protected switching fabrics
- Two protected controllers
- Protected power supply
- Up to three subracks in a standard ETSI or ANSI rack

Main interfaces of the Alcatel-Lucent 1850TSS 160 are:

- Data cards
  - 10 x Gigabit Ethernet (GE) packet module, Small Form- Factor Pluggable (SFP)
  - 10 GE packet module, 10 Gb/s Form-Factor Pluggable (XFP)
  - Multiservice packet over SONET/SDH (PoS) packet module, portless
  - ATM gateway packet module, portless
- SONET/SDH cards
  - 1 x OC-192/ STM-64: XFP
  - x OC-48/STM-16: SFP
  - 8 x OC-3-12/STM-1-4: SFP
  - 10 x any port card: a data/TDM concentrator: SFP
  - 1 x optical transport unit (OTU)-2: 10 Gb/s bidirectional transponder, tunable optics

The Alcatel-Lucent 1850 TSS is designed for low power consumption. Developing new components with very high integration density and low voltage supply leads to a significant reduction in power consumption. Basic components and related consumption are described in Table 7:

Table 7: Maximum power consumption for each board of Alcatel 1850 TSS-160.

Max power consumption for each board	
Board description	*Watts*
Shelf & Matrix incl. FAN and control platform (full protection)	266
1X10GE PACKET MODULE SYNCH ETH	100
10X1GE PACKET MODULE SYNCH ETH	100
MULTISERVICE PACKET MODULE	110
1X10G SYNC OPTICAL	40
4x2.5G SYNC OPTICAL	37
100BASE-FX SFP	1
100/1000BASE-T SFP	1.2
1000BASE-SX/LX/ZX SFP	0.85
STM-1/OC-3, STM-4/OC-12 SFP	1
STM-16/OC-48 I-16.1/SR-1, S-16.1/IR-1, L-16.1/LR-1 SFP	1.2
STM-16/OC-48, L-16.2/LR-2 SFP	1.5
10GBE BASE S XFP, I-64.1/10GE BASE L XFP	2.5
S64-2B/10GE BASE E XFP, DWDM XFP	3.5

4.1.4 IP routers

As in other kinds of network devices, operational power requirements of IP routers arise from all the HW elements realizing network-specific functionalities, like the ones related to data- and control-planes, as well as from elements devoted to auxiliary functionalities (e.g., air cooling, power supply, etc.). In this respect, the data-plane certainly represents the most energy-starving and critical element in the largest part of network device architectures, since it is generally composed by special purpose HW elements (packet processing engines, network interfaces, etc.) that have to perform per-packet forwarding operations at very high speeds. From a general point of view, IP routers have similar architectures with respect to high-end switching systems, since they usually have highly modular and hierarchical architectures.

For instance, the T-Series Juniper routers [30] have very complex and flexible architectures, which allow achieving 1.6 Tbps. These routers can be configured in multi-chassis setups (Figure 48) or in single ones. Each router chassis is substantially composed by a number of Port Interface Cards (PICs), Flexible PIC Concentrators (FPCs), a multi-plane switch fabric (Figure 49), and one or more Routing Engines. The PICs connect a T Series platform to the network and perform both physical and link-layer packet processing. They perform all of the functions that are required for the routing platform to receive packets from the network and transmit packets to the network. Many PICs, such as the Juniper IQ PICs, also perform packet processing. Each FPC may contain one or more Packet Forwarding Engines (PFEs). For instance, a T1600 100 Gbps slot supports two 50 Gbps PFEs. Logically, each PFE can be thought of as a highly integrated packet-processing engine using custom ASICs developed by Juniper Networks. These ASICs enable the router to achieve data forwarding rates that match fibre optic capacity. Such high forwarding rates are achieved by distributing packet-processing tasks across this set of highly integrated ASICs.

When a packet arrives from the network, the ingress PFE extracts the packet header, performs a routing table lookup and any packet filtering operations, and determines the egress PFE connected to the egress PIC. The ingress PFE forwards the packet across the switch fabric to the egress PFE. The egress PFE performs a second routing table lookup to determine the output PIC, and manages any egress class-of-service (CoS) and quality-of-service (QoS) specifications. Finally, the packet is forwarded to the network. The switch fabric provides connectivity between the PFEs. In a single-chassis system, the switch fabric provides connectivity among all of the PFEs residing in the same chassis. In a multi-chassis system, the switch fabric provides connectivity among all of the PFEs in the different chassis of the routing node cluster. In a single-chassis or a multi-chassis system, each PFE is considered to be logically contiguous to every other PFE connected to the switch fabric.

The Routing Engine executes the Junos OS and creates the routing tables that are downloaded into the lookup ASICs of each PFE. An internal Ethernet connects the Routing Engine to the other subsystems of a T Series platform. Each subsystem includes one or more embedded microprocessors for controlling and monitoring the custom ASICs, and these microprocessors are also connected to the internal Ethernet.

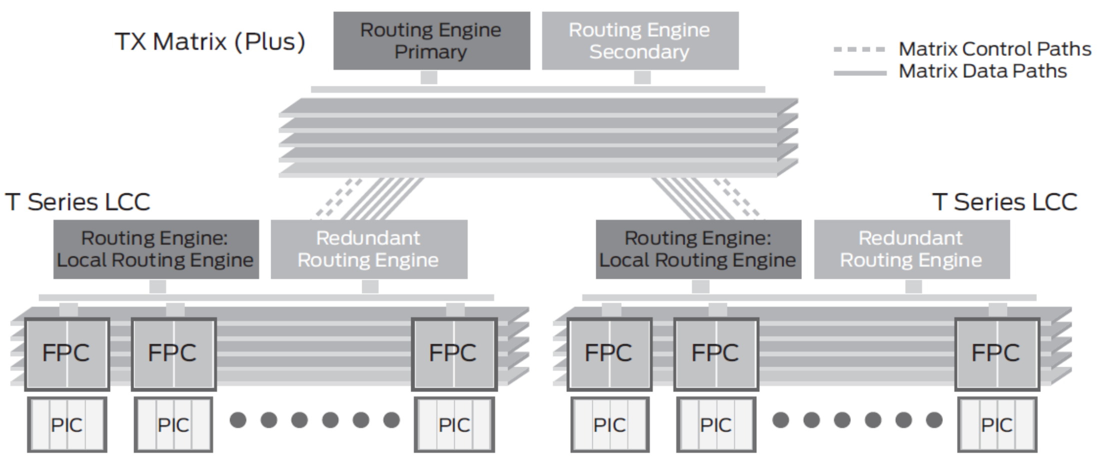

Figure 48: Multi-chassis system high-level overview of Juniper T-series.

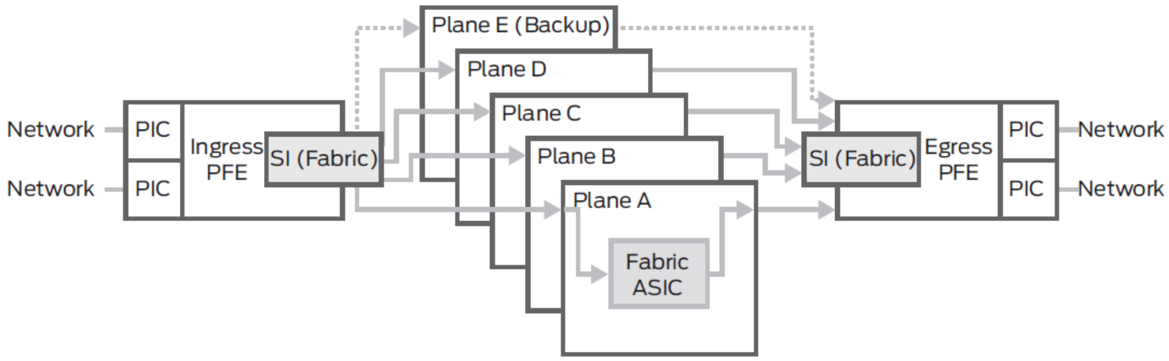

Figure 49: Five switch fabric planes for the T Series.

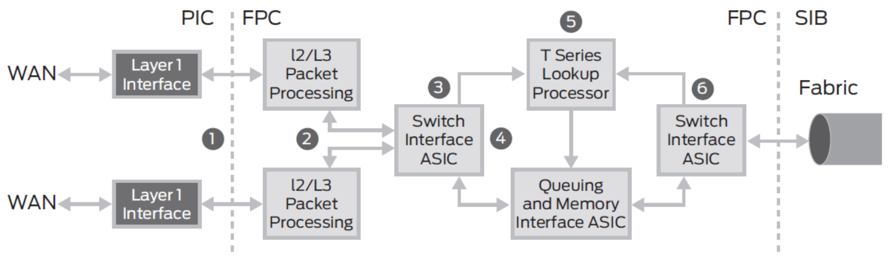

Figure 50: PFE and switch fabric using the T Series chipset ("to fabric" direction).

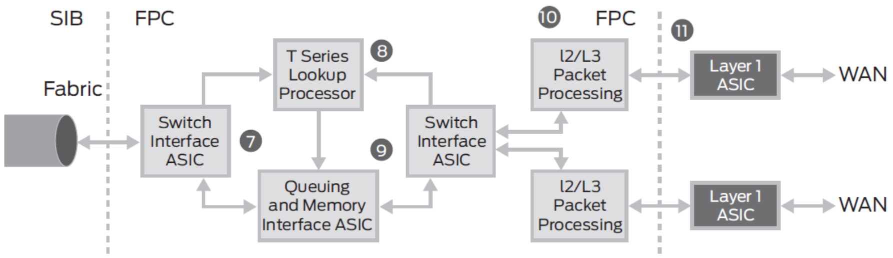

Figure 51: Packet flow from fabric to egress PFE.

As far as energy efficiency is concerned, the almost totality of IP router manufactures report that the energy consumption of their products is almost flat, and only depends on the number of HW elements/chassis/slot plugged in. For example, Juniper reported in [29] the energy requirements of some of their routers (see Table 8).

Table 8: Forwarding capacity and maximum power draw of different Juniper routers.

	M40	M160	T640	T1600
Slot Capacity	3 Gbps	10 Gbps	40 Gbps	100 Gbps
System Capacity	40 Gbps	160 Gbps	640 Gbps	1600 Gbps
Max System Draw	1.5 kW	3.15 kW	4.52 kW	8.35 kW

Till today, energy efficiency in IP routers substantially means only maximizing the ratio between the system capacity and its power consumption, without taking the possibility of adapting system performance to the real workload into account. For this reason, internal router components do not include power management capabilities like standby, low power idle or performance scaling.

Tucker et al. [18] and Neilson [27] focused on high-end IP routers, and estimated that the data-plane weighs for 54% on the overall device architectures, vs. 11% for the control plane and 35% for power and heat management (see Figure 52). The same authors further broke out energy consumption sources at the data-plane on a per-functionality basis. Internal packet processing engines require about 60% of the power consumption at the data-plane of a high-end router, network interfaces weigh for 13%, switching fabric for 18.5% and buffer management for 8.5%.

Notwithstanding that this study specifically refers to high-end router platforms, and the same internal distribution of power cannot be obviously maintained for all the typologies and the architectures of network devices, the resulting estimations provide a relevant and clear indication on how and where future research efforts need to be focused in order to build next-generation green devices.

Following these basic ideas, the largest part of current research contributions generally focused on introducing novel extensions and solutions for reducing the carbon footprint of particular devices, by working at the level of internal processing engines for core and transport network nodes, and that of network interfaces and/or link protocols for access and home equipment.

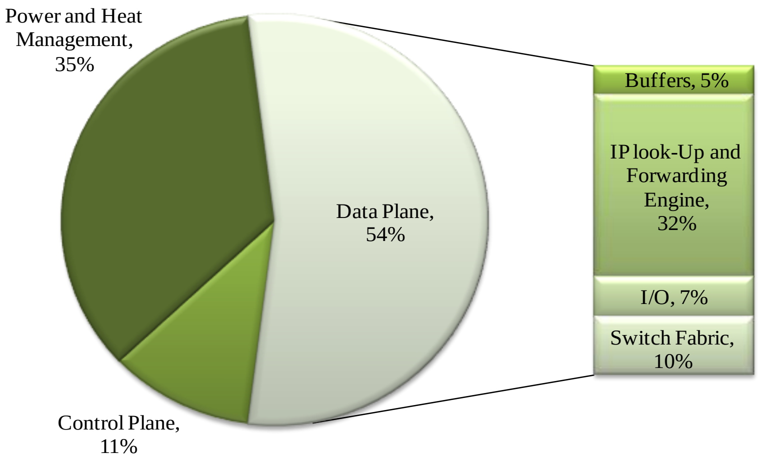

Figure 52: Estimate of power consumption sources in a generic platform of high-end IP router. Source: [118].

4.1.5 Network cards

In this section, we report some measurements performed on off-the-shelf Gigabit Ethernet adapters for PC platforms. In more detail, we measured the energy consumption of 3 adapters from Intel, namely the quad- and dual-port versions of the Intel PRO-1000-ET base Tx, and the dual port version of the Intel PRO-1000-PT base Tx.

The obtained results show that the energy consumption does not depend on the real utilization of the boards, but on the number of attached cables and the speed of links (see Figure 53-Figure 56). The overall consumption ranges between about 4 W to 13 W for all the considered adapters. Strangely enough, we observed that the power consumption of the quad-port adapter is lower than the dual-port ones, and that the adapters with attached cables running 10 Mbps speeds waste less energy than the same adapters without attached cables. This is a clear indication on the lack of energy-optimization in the design of such boards.

Figure 53, Figure 55 and Figure 56 exhibit that the adapter energy consumption directly depends on the number and the speeds of Ethernet links: e.g. using the 100 Mbps speed allows reducing the energy absorption of 30-50%.

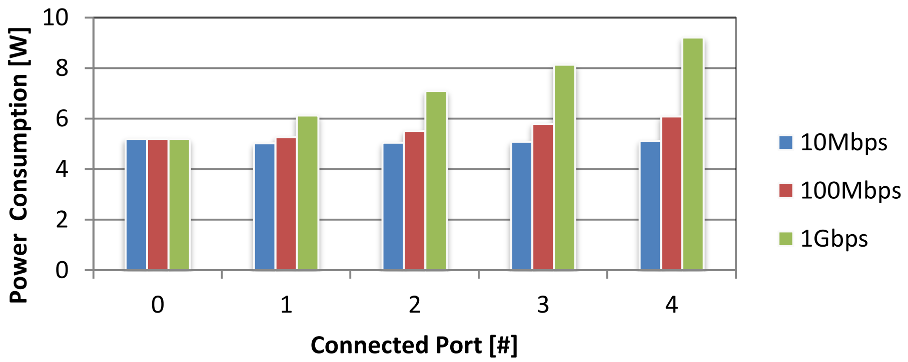

Figure 53: Power consumption for the Intel Gigabit ET quad-port port with 1000-BaseTx ports, and with respect to the number of cables physically attached and the speed of Gigabit links.

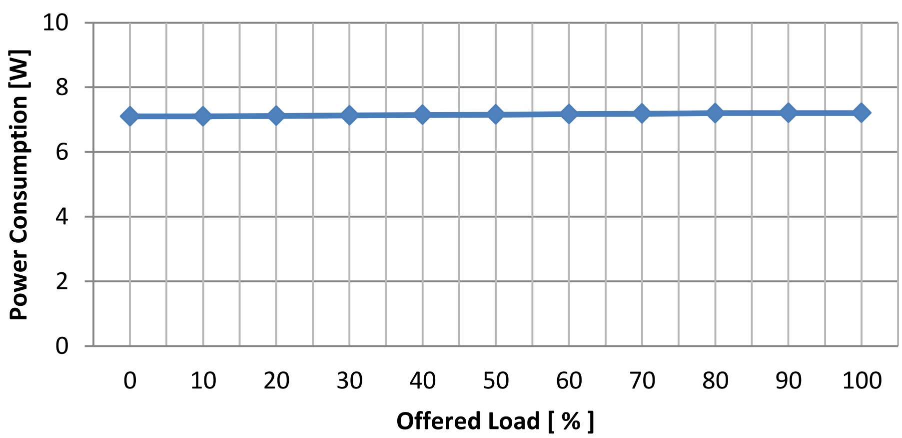

Figure 54: Power consumption for the Intel Gigabit ET quad-port port with all the 1000-BaseTx ports enabled and running at the Gigabit speed according to the traffic offered load.

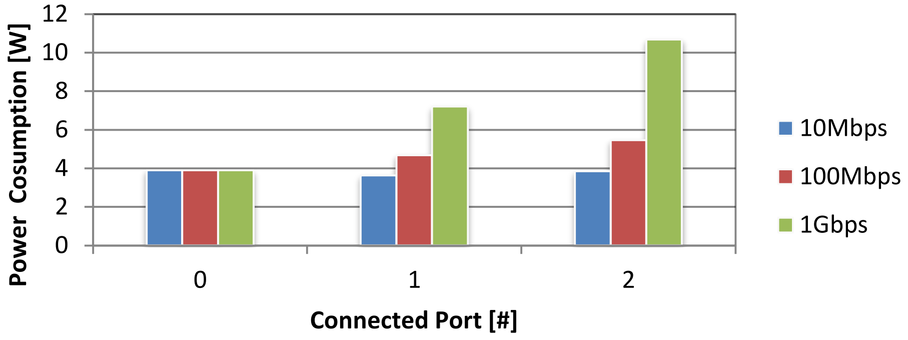

Figure 55: Power consumption for the Intel Gigabit ETdual-port port with 1000-BaseTx ports, and with respect to the number of cables physically attached and the speed of Gigabit links.

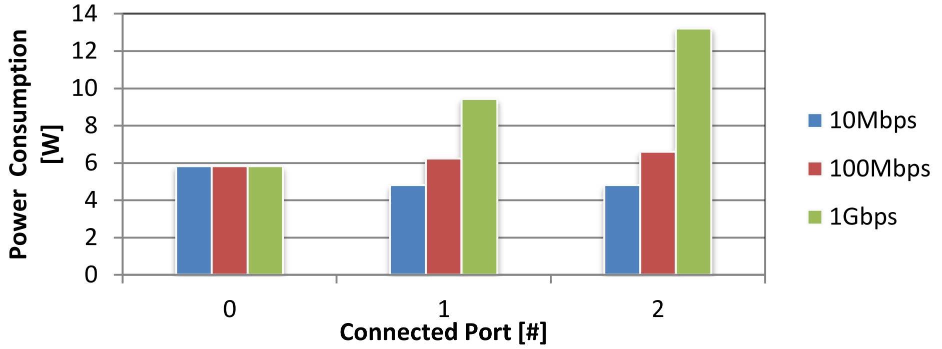

Figure 56: Power consumption for the Intel Gigabit PT dual-port port with 1000-BaseTx ports, and with respect to the number of cables physically attached and the speed of Gigabit links.

Moving to higher speeds, the power consumption of a Host Channel Adapter (HCA) card is mostly composed of the power used by the network processor on the card itself. Most of the interface cards are cooled using passive cooling that does not consume more power. Most of network cables used to be passive copper based, hence consuming no power. However, nowadays there is a growth in use of optic active cables that support bandwidth incensement towards 10/40Gbit Ethernet and 56/80 Gb/s InfiniBand protocols and support longer distance connectivity.

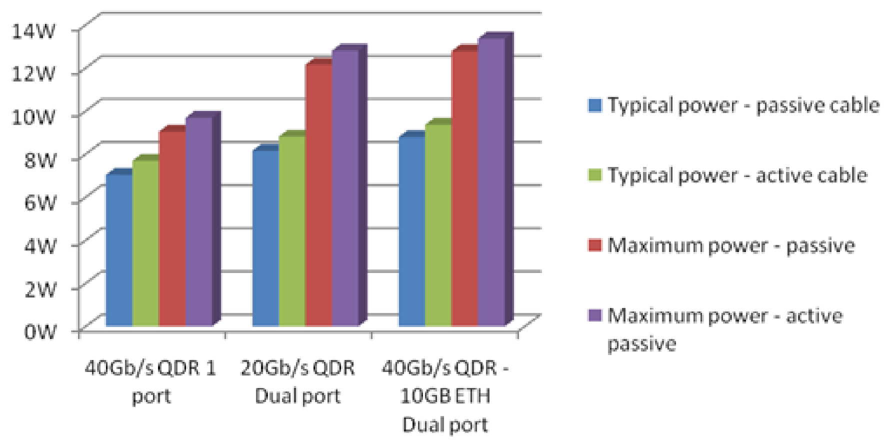

Figure 57: Power Consumption of Mellanox Connect-X2 Family.

Figure 57 above displays the typical and maximum power consumption of Mellanox Connect®-X family. The numbers on the left side of Figure 57 reflect the power consumption of the network interface and the cables that are connected to it. When discussing the energy consumption of network interfaces, we should consider the hardware on which the network interface is mounted, generally, a server or a data storage unit. A server's power budget can be significant larger and measured in hundreds of watts. Thus it is prominent to discuss Wake on Lan (WOL) technologies when considering power saving techniques in the network end points. In WOL mode the power consumption is between 1W - 3W and is highly depended on the network technology (e.g. 1/10/40Gb/s) and the type of cable used (active / passive).

4.1.6 Direct measurements in off-the-shelf devices

In this subsection, specific results are presented from experiments that were performed to determine how much energy is consumed by typical network devices under various load conditions as well as when they are in the idle state.

Direct measurements on Cisco 1700 series routers

In the first experiment, power consumed by four Cisco 1720 routers was measured jointly in order to make the constant component of measurement error less significant. The devices are connected by means of Serial and Fast Ethernet interfaces, as shown in Figure 58. There is total of 5 serial interface cards in the system, two of which are WIC T2 with two serial interfaces on-board, and the other three cards are WIC T1. Four serial interfaces are connected, while the rest is down during the experiment. The clocks of serial links were set to 8 Mb/s while the network was loaded with TCP traffic to maximum capacity of Interfaces. The current was measured with a clamp-on current meter. To test the control plane influence on power consumption, traffic shaping, policing, and class-based weighted fair queuing (CB-WFQ) were turned on. Traffic was generated accordingly, with several destination port numbers exploited, to make the packets end up in different class-based queues.

The results from the measurements on Cisco 1720 routers are shown in Table 9.

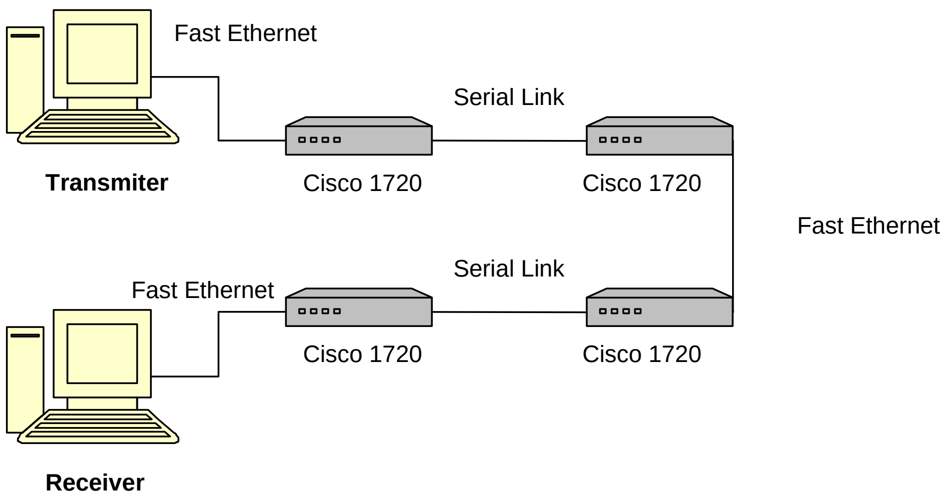

Figure 58: Testbed topology.

Table 9: Measurements on Cisco 1720 routers.

	Conditions	Measured current @ 230V mains cord [A]
1.	Routers fully loaded. Policing, shaping and class based queuing turned on.	0.2
2.	Routers fully loaded. Policing, shaping and class based queuing turned off.	0.2
3.	No traffic	0.2
4.	Serial interfaces in state of administrative shutdown	0.19
5.	Serial cables disconnected	0.18
6.	WIC cards physically removed	0.13

In the second run the results were verified by the Kyoritsu multimeter connected in series, as it is shown in Table 16.

Table 10: Measurements on energy consumption of Cisco 1720 routers made by multimeter.

Conditions	Measured current @ 230V mains cord. [A]
Routers fully loaded. Policing, shaping and class based queuing turned on.	0.175
No traffic	0.170

As it can be noticed, traffic load has very little, or no impact on power dissipation in Cisco 1700 series routers. However, physical removal of the serial interface card seems to have great influence on power consumption. This leads to the obvious conclusion that the interface cards should be re-designed so that power consumption is reduced in idle periods.

Direct measurements on Cisco Catalyst 2900 XL series switch

In the second experiment, the Catalyst 2900 switch is examined that is a twenty-four port Fast Ethernet switch that enables VLAN configuration. In order to measure energy consumed by several interfaces working at the same time, three VLANs were defined, and connected in-series (stacked) as shown in Figure 59. The links were saturated with TCP traffic to maximum capacity of interfaces, i.e. 100Mb/s. The achieved maximum transmission speed during the experiment was about 94 Mb/s. All measurements were performed with Voltcraft Energy Logger 4000.

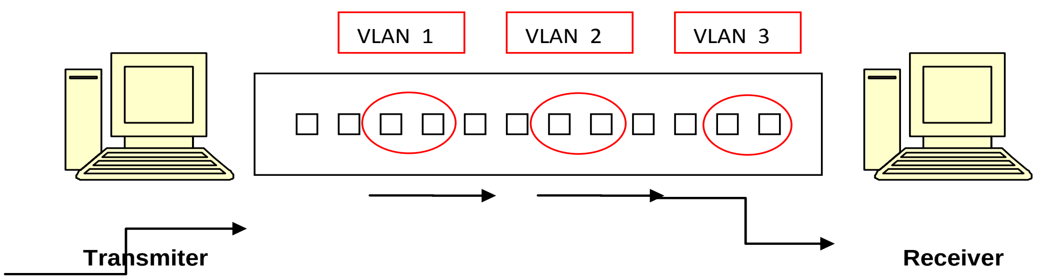

Figure 59: Wiring in testbed for Experiment #2.

Table 11: Measurements on Cisco Catalyst 2900 XL series switch.

	Conditions	Power [W]	Apparent power [VA]	cos φ	Voltage [V]	Current [mA]
1.	All Interfaces UP, but with no wires connected.	39.7	65.2	0.6	220.9	0.29
2.	Two interfaces wired to hosts. All ports in administrative UP state. No traffic	40.1	66.2	0.6	220.9	0.3
3.	Two interfaces connected, wired to hosts, in UP state. Loaded with “bulk” TCP traffic generator	40.1	66	0.6	221	0.29
4.	3 VLANs configured, 6 interfaces UP, 18 interfaces administratively down.	32.2	53.9	0.6	221	0.24
5.	3 VLANs configured, 6 interfaces UP, 18 interfaces administratively down, 6 cables inserted.	33.3	54.9	0.6	221	0.24
6.	3 VLANs configured, 6 interfaces UP, 18 interfaces administratively down, 6 cables inserted. Loaded with “bulk” TCP traffic generator	33.6	56	0.6	221	0.24
7.	3 VLANs configured. All 24 interfaces up, 6 cables inserted. Loaded with “bulk” TCP traffic generator.	40.4	66.1	0.6	221.9	0.29
8.	All interfaces Administratively down	29.2	49.6	0.59	221	0.22

Based on the measurements’ results shown in Table 11, it can be claimed that traffic load has very little impact on power dissipation. Very interesting observation is the existence of low power consumption when all interfaces are in “administratively down” state.

Direct measurements on Cisco catalyst 6506-E switch

In the third experiment, energy consumption on the Cisco catalyst 6506-E switch is examined. The switch was equipped with a single power supply, two processor cards, two interface cards and a set of fans (Table 12). All power measurements were performed with Voltcraft Energy Logger. Layer 2 traffic was generated by means of JDSU HST3000 and JDSU Smart Ethernet pair of probes.

Table 12: Cisco catalyst 6506-E switch cards.

Slot	Details
G1	16 SFP ports
G2	48 1000 base T ports
G5	Processor, 8 x GBIC ports, 1 x 1000 Base T
G6	Processor, 8 x GBIC ports, 1 x 1000 Base T

Table 13: Measurements on Cisco catalyst 6506-E switch.

	Conditions	Power [W]	Apparent power [VA]	cos φ	Voltage [V]	Current [A]
1.	All interfaces are in “administratively down” state	509.7	532.8	0.95	227	2.34
2.	All interfaces on card in slot G1 up	508.8	532	0.95	227	2.34
3.	All interfaces on cards in slot G1 up + G2 up	507.6	530.7	0.95	227	2.33
4.	3 VLANs configured in G2 card. No traffic	507.7	530.5	0.95	227	2.33
5.	As above + 6 cables plugged in	512.3	535.1	0.95	226.2	2.36
6.	1 Gb/s UDP traffic going through 3 VLANs connected in series.	522.4	544.9	0.95	226	2.41
7.	1 Gb/s UDP traffic going through 2 VLANs connected in series.	517.6	540	0.95	226	2.38
8.	1 Gb/s UDP traffic going through 1 VLAN.	512.6	534.9	0.95	226	2.36
9.	100 Mb/s UDP traffic going through 3 VLANs connected in series.	513.8	536	0.95	226	2.37
10.	Chassis only without cards	128.5	161.8	0.79	227.3	0.71
11.	One of processor card remaining, the rest removed.	238.6	269	0.88	227.3	1.18
12.	Ethernet card in slot G2 + 1 processor	311.8	338	0.91	227	1.48
13.	Cards in slots G1 + G2 + 1 processor	397.9	423.7	0.93	227	1.86

As it is shown in Table 13, unlike in previous experiments, shutting down inactive ports (without wires attached) does not reduce power consumption. Possible explanation is that the cards already do implement some power saving mechanisms on interface level. Any traffic-related configuration activities, like connecting wires, putting interfaces in UP state and setting up VLANs result in significantly greater power consumption. Power consumption in typical configurations is also presented graphically in Figure 60. Although the power consumption adapts well to traffic load, still the cost of ports that are switched off accounts to nearly 40 per cent of the costs of ports fully loaded.

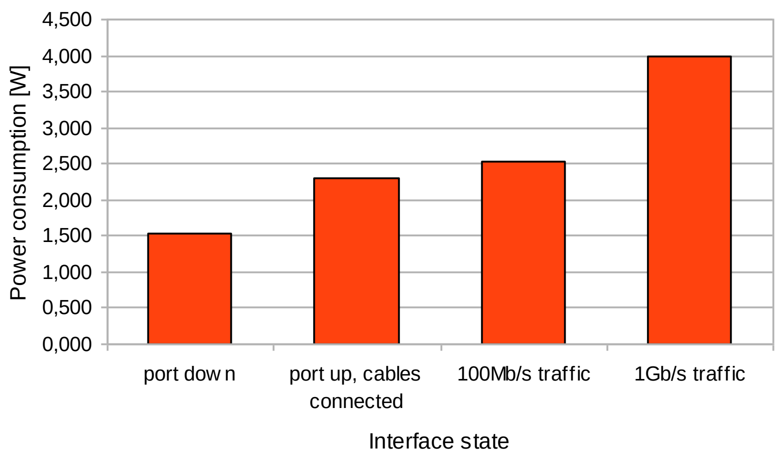

Figure 60: Power consumption per one 1000-BaseT port of 48-port Ethernet card.

Direct measurements on Cisco 7201 router

In the fourth experiment, power consumption on Cisco 7201 router is examined. In this experiment, layer 3 traffic was generated by means of JDSU HST3000 and JDSU Smart Ethernet pair of probes.

Table 14: Measurements on Cisco 7201 router.

	Conditions	Power [W]	Apparent power [VA]	cos φ	Voltage [V]	Current [mA]
1.	Two interfaces wired to hosts. No traffic	54.6	61.9	0.87	218.7	0.28
2.	1 Gb/s UDP traffic, one way, packet size 1500B.	55.5	62.5	0.88	218.6	0.28
3.	100 Mb/s UDP traffic, one way, packet size 1500B.	54.7	61.9	0.88	218.7	0.28
4.	1 Gb/s offered UDP traffic, one way, packet size 64 B. 550 Mb/s passed through.	57	63.8	0.89	219.3	0.29
5.	100 Mb/s UDP traffic, one way, packet size 64B.	54.6	62.3	0.88	219	0.28
6.	1 Gb/s UDP traffic, two way, packet size 1500 B.	56.7	63.4	0.89	218.7	0.29
7.	100 Mb/s UDP traffic, two way, packet size 1500 B.	54.7	62.2	0.88	219	0.28
8.	1 Gb/s offered UDP traffic, two way, packet size 64 B. About 300 Mb/s passed through.	57.1	64	0.88	218	0.29
9.	100 Mb/s UDP traffic, two way, packet size 64 B.	54.7	62	0.88	219.0	0.28

According to the results shown in Table 14, it is shown that the small number of interfaces limits the precision in measuring their power consumption. However, the rise of consumed power when switched from 1500 B packets to 64 B packets, which resulted in approximately a 7-fold rise in number of packets served, suggests that the data plane contributes significantly to overall power consumption of the router. This observation should be taken into account while developing distributed control strategies as not only the traffic volume in kB/s, but also in packets/s, should be taken into account while making routing decisions.

4.2 Reference network scenarios

In this subsection, the three reference scenarios that will be examined within ECONET are presented in detail. The first scenario regards a telecom operator network while the upcoming two scenarios regard Internet Service Providers networks. In each case, the network topology, the traffic profiles and volumes are presented and the energy-aware techniques that are going to be studied within ECONET are analysed.

4.2.1 Reference Telecom Operator scenario by TELIT

Network topology

The Telecom Italy Wireline IP/MPLS platform is composed of a national backbone OPB (Optical Packet Backbone) conceived in early 2001 as an evolution of a previous IP network called IBS (InterBusiness) to bring together all voice and data services on a single network platform. The OPB network is modelled on 32 PoPs to concentrate all data access technologies currently in the field in the main national transmission nodes.

This architecture, purged from the ATM transport component, and with innovative features and performance, has enabled several important goals. One such goal is, for example, the provision of a geographical transport service among the Media Gateway's BBN (national backbone) of dedicated telephone traffic over IP, with guarantees of high protection for transmission failures, priorities of treatment packages during the phenomena of Network congestion and safety compared to the rest of the data traffic is being carried. This achievement has allowed the gradual demise of the traditional national long-distance telephone network and the integration of the transport of IP data traffic with telephony traffic on the OPB platform.

As far as the network design criteria are concerned, the entire network, except link terminations, is specifically dimensioned to have maximum peak utilization equal to 50% of allowable bandwidth in each link at every level. It is worth noting that 50% is only the maximum allowable link utilization, and real utilizations are often much lower than this figure.

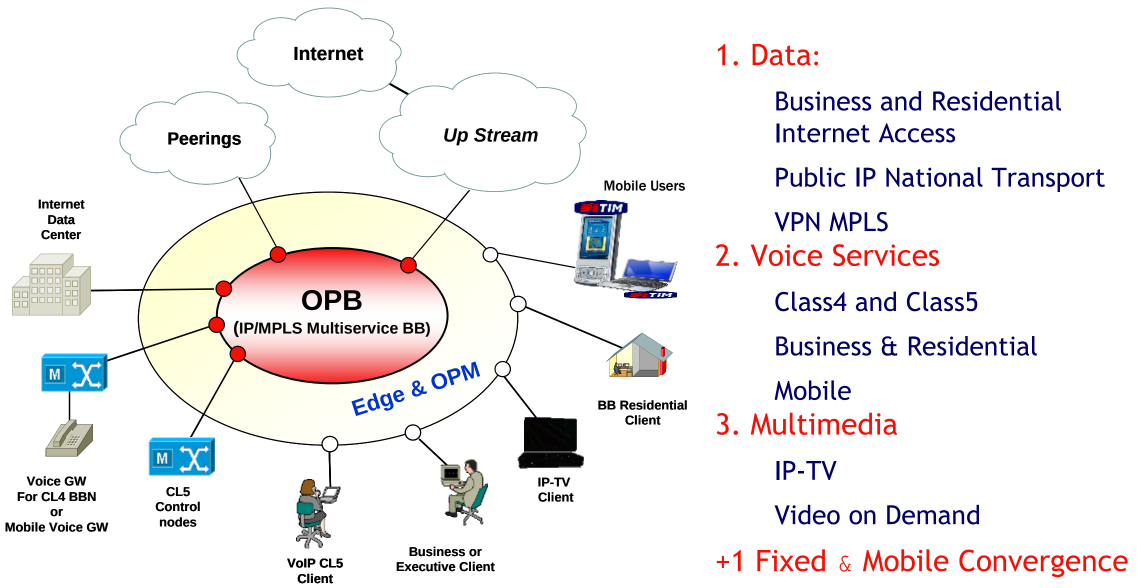

Figure 61: Services on the two backbone networks.

Figure 61 summarizes an example of the possible breakdown of services offered and transported through the backbone. The OPB network is directly connected to the networks of other ISPs (Internet Service Provider) through "Private" and "Public" peering points, as appropriate, generally established in one of the two Italian neutral interconnection points in Rome and Milan. The interconnection between the two backbones is done in Rome and Milan. Connectivity to the Big Internet, for all destinations that are not internal or covered by Peerings just mentioned, is obtained through the International backbone of Telecom Italy. The entire national clientele, with or without its own Autonomous System, for the Full Internet or MPLS VPN services, access the backbone via appropriate Edge structures (Access Router, MPLS PE routers, NAS and NAS ADSL Dial-up) by using the many dedicated access networks ATM, ADSL, GBE, SDH, and ISDN.

OPB network architecture

The backbone OPB (Optical Packet Backbone) is divided into different national PoPs, with a double-centre star topology and the distinction between Inner Core and Outer Core PoPs (Figure 62). The Inner Core is made up of two PoPs in Rome and two PoPs in Milan, while the Outer Core is focused on others peripheral PoPs connected in "dual-homing" to Rome or to Milan on the basis of their geographical location. The “dual-homing" of the Outer Core PoPs to the Star Center was made even more reliable doubling the PoP in Rome and Milan on different central offices and connecting each of the Outer Core PoP with at least 2 transmission circuits to the two Inner Core PoPs.

At the transmission level the OPB nodes are interconnected through the National Transport Optical Network. This network is fully redundant in all its components and it is immune from single fault conditions, such as the failure of a router, of a single card or of a transmission circuit. In the case of two simultaneous failures, it is up to the functionality of QoS / CoS (Quality of Service, Class of Service) to safeguard the most valuable services. The project sizing of the circuits imposes a limit on the maximum total traffic for each link equal to 50% in order to be immune to a single fault condition.

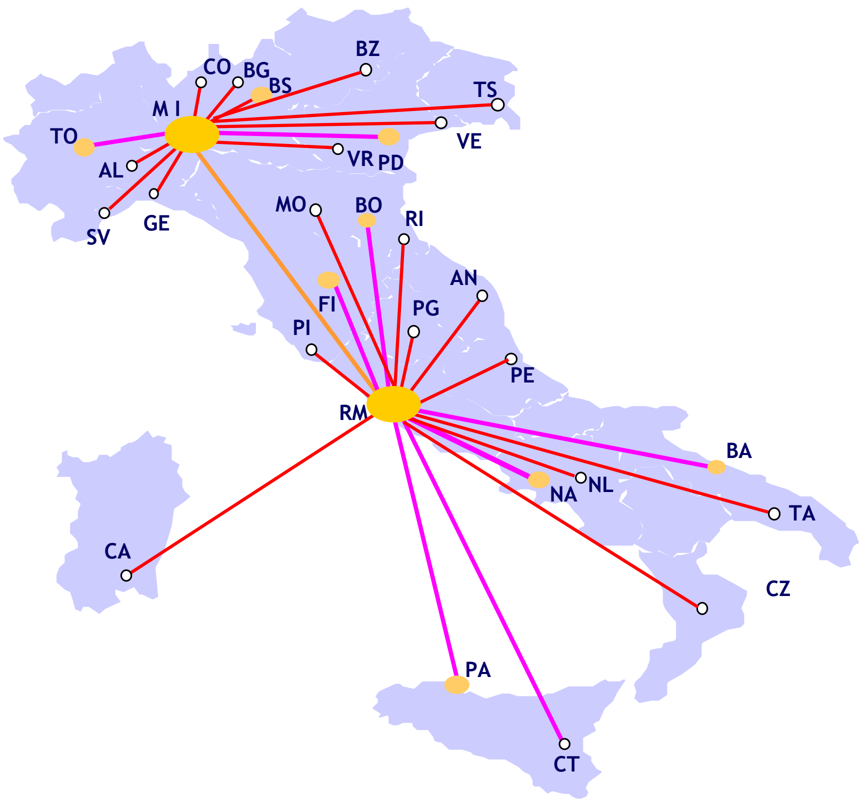

Figure 62: Backbone OPB architecture.

The Inner Core

The network Inner Core is divided into the 4 PoPs in Rome and in Milan, independent and self-consistent (Figure 63), each with high-capacity traffic router equipment that covers the role of:

- Star Center (CS) for the circuits termination to the Outer Core PoPs;
- Transit (T) to the structures of Edge / Access;
- International Gateway (ITZ) for the termination of the connections to the Seabone PoPs;
- Gateway (P) to the national peering points.

The four PoPs are based on a double quadrilateral of transmission circuits for the four backbone connections Rome-Milan. Moreover they are active redundant metropolitan links in order to connect the four Star Centres to the International and Peering Gateways and to balance the traffic into the network. All PoPs of the Outer Core are terminated to a pair of Inner Core devices with the function of Star Centres that are physically located in two separate Inner Core POP's. To further increase the reliability, each Inner Core PoP is fully redundant and the links from Outer Core PoPs are distributed on two pairs of nodes. The four Inner Core PoPs, besides collecting and distributing traffic throughout the network, have to aggregate traffic and services locally for Rome and Milan, using the Edge structures devoted to the services Executive / Business or Residential.

Figure 63 shows the block diagram of the different network layers.

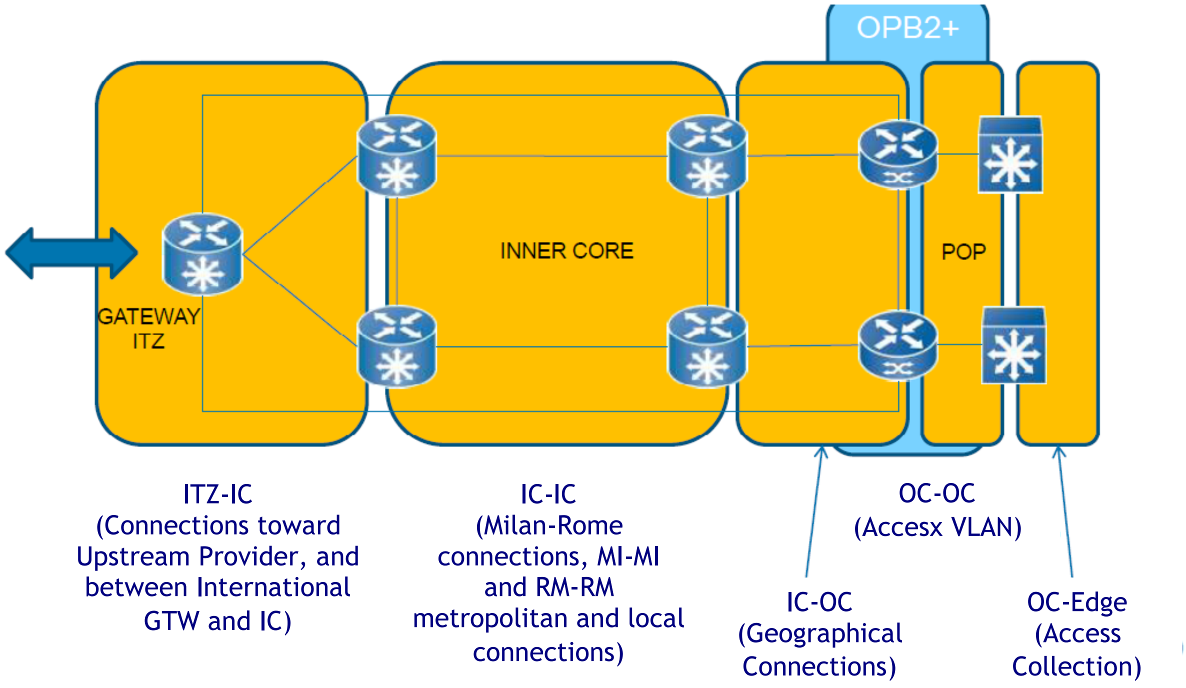

Figure 63: Network layers' scheme.

The Outer Core

The Outer Core is made up of 28 national PoPs, main or secondary depending on whether they are sites with more or less traffic importance. From the functional point of view both types of Outer Core Outer PoPs have the same characteristics and what differentiates them is the quantity of traffic and the number of devices involved.

The architecture of a generic network PoP is shown briefly in Figure 63: there is always a pair of core routers, called Transits, connected in redundant configuration to the Inner Core, and a pair of aggregation devices that collect all access systems/devices (Provider Edge Router, BNAS, etc.). Typically, the "access" structures in this context correspond to level 3 equipment (that is IP routing and not level 2 switching) through which final customers are served or are provided Service / Control components. To this purpose, a pair of switch devices with level 3 functionality creates two access VLANs, on load sharing and fully redundant, for the following "Access" structures:

- Edge Executive, through the Access Router (AR), for Full Internet Services, and Router PE (Provider Edge) for other MPLS services;
- NAS ADSL for residential or small Business customers;
- NAS dial-up.

The traffic allocation is made by balancing the IP / MPLS packets over four different routes, among the Access routers and the first two OPB routers, on the basis of the "flow" by using a hashing algorithm that takes into account the source address, and destination TCP / UDP ports. The Voice Gateway of the BBN Network (transit nodes of the telephone network) and the front-end to the network of metropolitan traffic collection (network OPM - Optical Packet Metro) are, instead, directly connected to the transit routers.

Just as an example, Figure 64 reports the network segments involved for the provision of Internet access service, both for fixed and mobile access:

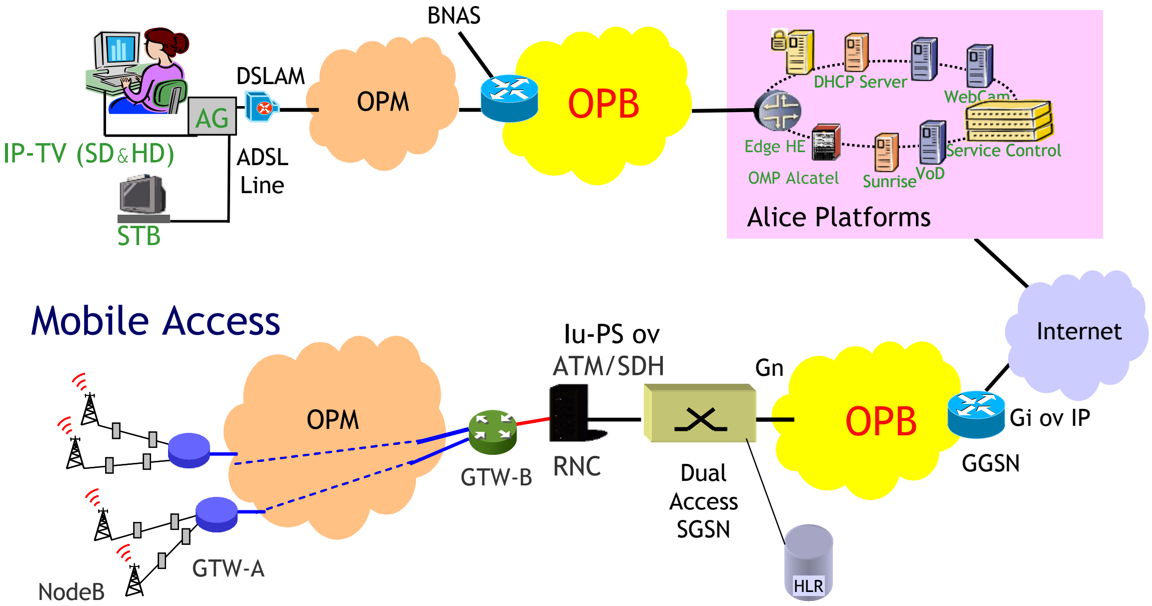

Figure 64: Example of structure for Internet access services provisioning.

The Access Network

Telecom Italia aims at deploying in the next years its Next Generation Access Network (NGAN) as a mix of different architectures and technologies, as shown in Figure 65.

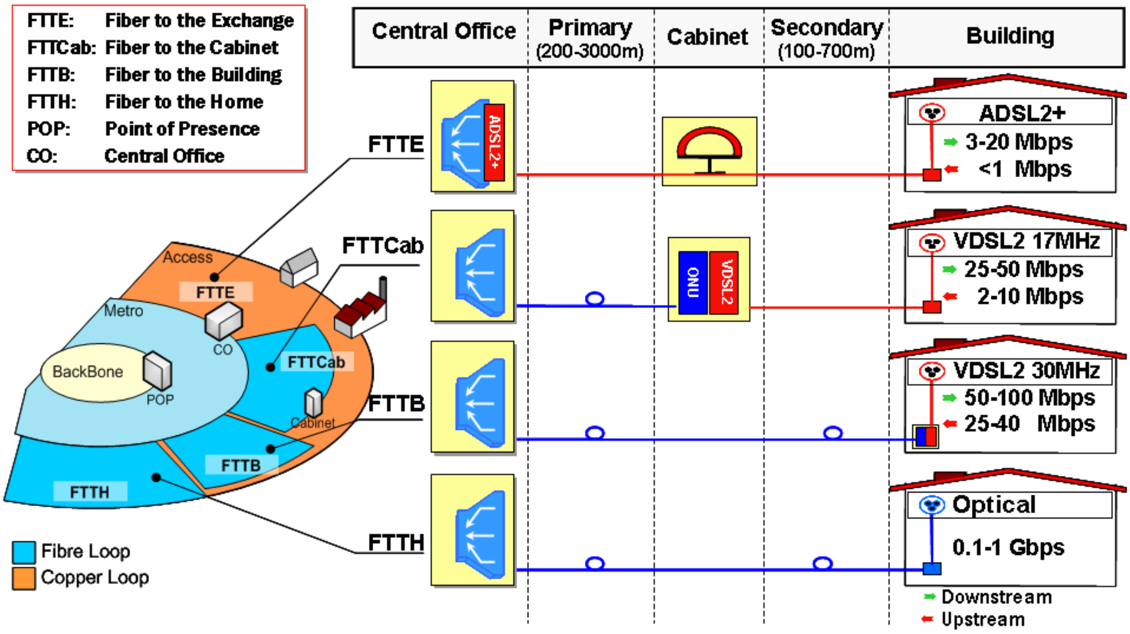

Figure 65: the NGAN architecture of Telecom Italia.

The current architecture used to provide the Broadband service is Fibre-To-The-Exchange (FTTE); this means that the fibre "stops" at the Central Office (CO), and the connection from the CO to the customer's home is still provided with the traditional copper pair using the ADSL2+ technology. Even if the access network of Telecom Italia is one of the world's shortest (with an average of 1.8km and 99% of the lines shorter than 5.5km), with the FTTE architecture the downlink throughput cannot go over 20Mbit/s. Therefore, in order to increase the speed it is needed to shorten the loop, deploying the fibre closer to the customer. This can be done, increasingly, with:

- Fibre-To-The-Cabinet (FTTCab), where the fibre stops at the Street Cabinet and the copper-based secondary network is re-used (reducing the loop to 100÷700 m). This can enable downlink speeds up to 50Mbit/s, through the use of the VDSL2 technology (with profile 17a).
- Fibre-To-The-Building (FTTB), where the fibre reaches the basement (or the very neighbourhood) of the building and the copper drop is therefore reduced to less than 150 m. In this case, the speed increases up to 100Mbit/s, thanks to the use of the VDSL2 technology (with profile 30a). This option is currently taken into account only as an alternative to FTTH, for places where this can hardly be implemented.
- Fibre-To-The-Home (FTTH), where the fibre is directly brought to the customer's apartment. This is of course the fastest architecture, with speeds in the range 100 Mbit/s ÷ 1 Gbit/s.

Provided that the best scenario, from both the throughput and the energy consumption point of view, would be the FTTH, it must reasonably be considered that from an economic point of view it would be absolutely not feasible. That's why Telecom Italia will deploy all the three architectures described above according time by time to geographical, infrastructural and profitability aspects.

Focusing now on FTTH, it must be highlighted that it can be realized using two different solutions: the Point-to-Multipoint (that is currently represented by the GPON technology) and the Point-to-Point (that is currently represented by the Ethernet technology). In the first case, the basic principle is that a single optical fibre is "shared" among multiple users, while in the second case each customer gets a dedicated fibre. It is commonly known that the Point-to-Multipoint solution is typically preferred by the Incumbent Operators, while on the other hand the Point-to-Point solution is the main choice of the OLOs: this because it replicates the current paradigm of the traditional copper network, allowing the OLOs to keep their current processes, and also because they can limit their investments in areas where they have a very low market share. Coming back to the GPON solution, it is basically formed by:

- an active equipment (Optical Line Termination – OLT), installed in the Central Office with typically 64-160 GPON ports per shelf
- an active equipment (Optical Network Termination – ONT), which terminates the optical line in the customer's home
- an Optical Distribution Network (ODN), deployed between the Central Office and the customers, which is completely passive and is constituted by the optical fibre and a set of optical splitters. These splitters allow the partition of the input signal into different output signals (typically from 4 to 16) and vice versa. The portion of the ODF served by a single GPON port of the OLT is called a GPON tree.

The current GPON solutions can allow an overall splitting factor of 1:128, which means that up to 128 customers can share the same optical fibre and therefore also the associated total bandwidth, which for the ITU-T G.984.x GPON standard is roughly 2.5Gbit/s in downlink and 1.25Gbit/s in uplink. The maximum distance from the CO to the customer allowed by the GPON technology is 20km. The current FTTH GPON architecture for Telecom Italia has been reported in Figure 66: as visible, two splitting levels (one at the street manhole, with ratio 1:N and the other in the building basement, with ratio 1:M) are foreseen. In case of an overall splitting ratio of 1:128, typically 90 users are actually connected, while in case of splitting ratio of 1:64, the average number decreases to 50.

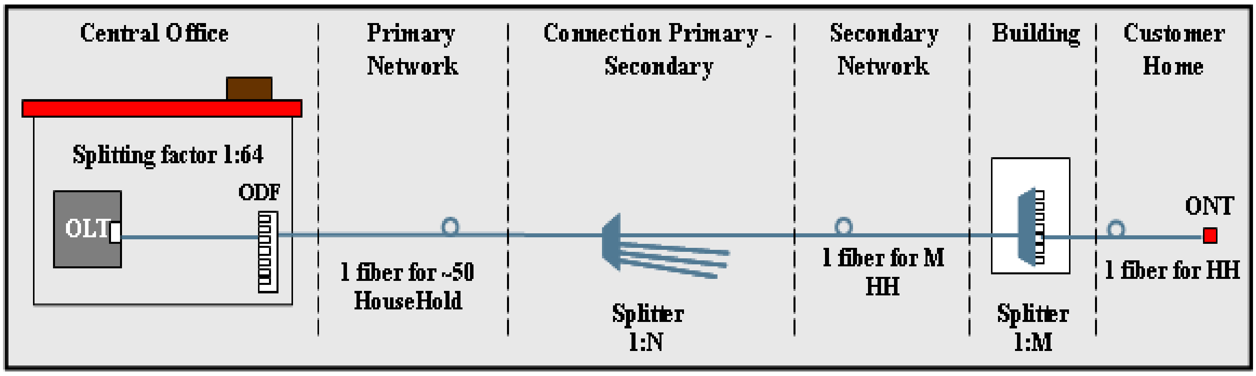

Figure 66: FTTH GPON architecture for Telecom Italia.

The GPON solution is absolutely future proof and can therefore guarantee the investment performed on the infrastructure: as a matter of fact, the Standardization Bodies FSAN and ITU have already defined an evolutional path that will guarantee a progressive increase in the bandwidth and in the symmetry as well. In particular, during 2011 the new technology XG-PON1 (providing 10Gbit/s in downstream and 2.5Gbit/s in upstream for each GPON tree) will be already commercially available. Moreover, this new technology will allow the coexistence with the old GPON systems.

TELIT network equipment and typical energy consumption

The current situation and the historical data regarding the energy requirements of the TELIT network infrastructure were already introduced in section 3. The objective of this section is to make a step beyond these data and to evaluate how these energy figures will change in the next-future. To this purpose we worked on one of the most probable evolutions of the TELIT network, which would be deployed by 2015-2020.

We refer to a network with 17.5 million customers, where we assume the presence of broadband-only access technologies, together with suitable over-provisioning in the Metro, Transport and Core segments. Moreover, the analysis proposed here, aims at considering not only direct energy consumption of TELIT equipment, but also the one “induced” to end-costumers by considering network terminations. The considered reference scenario is shown in Table 15, which contains the number of devices per network segment and their energy consumption, in the business-as-usual case (i.e., the case in which no green enhancements would be included in network devices). Both end-user and operator equipment have been taken into account.

Devices’ energy consumption has been forecasted on the basis of present values, high quality specifications (e.g., European Broadband Code of Conduct), and expected “inertial” technological improvements (e.g., Dennard’s law [2]). In the devices’ energy requirements, we included the contribution of site cooling and powering systems, which account for 36% of device direct consumption. Starting from the scenario in Table 15, the per-user average energy requirement consists of about 111 kWh per year, and it is mainly due to home and access networks, for 79% and 16%, respectively. Metro/transport and core networks account only for 5%, but their joint energy requirement of about 107 GWh per year (about a quarter of the Telco’s direct energy consumption) can be a convincing driver for reducing the carbon footprint of backbone devices in the near future.

Table 16 defines key parameters that allow synthetically representing the average usage of network devices and links. This has been done by defining the expected (by 2015-2020) average customer up-times and loads, the average traffic utilization on metro/transport and core networks, and the number of devices and links that are usually deployed for redundancy purposes. It is worth noting that the traffic load values in Table 16 are significantly larger (and, consequently, give rise to conservative consumption estimations) than those of the current network and what is indicated in other studies [2].

Table 15: 2015-2020 network forecast: device density and energy requirements in the business-as-usual case (BAU). Example based on the Italian network.

	power consumption [W]	number of devices [#]	overall consumption [GWh/year]
Home	10	17,500,000	1,533
Access	1,280	27,344	307
Metro/Transport	6,000	1,750	92
Core	10,000	175	15
Overall network consumption			**1,947**

Table 16: Traffic and Topological data (average figures) for the 2015-2020 perspective network. Source: Telecom Italia.

Home access	number of customers per DSLAM	640
	usage of a network access (user up-time)	30%
	link utilization when a user is connected	10%
Core, transport and metro	redundancy degree for metro/transport devices (η_t)	13%
	redundancy degree for core devices (η_c)	100%
	redundancy degree of metro/transport links (ψ_t)	100%
	redundancy degree of core device links (ψ_c)	50%
	link utilization in metro networks	40%
	link utilization in core networks	40%

Traffic profile for the access network

The traffic volume and the relevant profiles depend on several parameters that go from the characteristics of the served area to the time of the day and the type of day (working or holiday, for example). In the following paragraphs we show typical profiles in different conditions for the links between the edge router and the DSLAM. The traffic in this part of the network is heavily affected by the final customers' activities due to the fact that is the last segment before the final connection to the customer termination (link between DSLAM and Access Gateway or other type of network terminations).

The curves shown represent four different conditions; two of them represent the traffic profile typical for a link through an area with prevalent business clients, one in a working day and the other in a holiday; the other two are typically for links to a DSLAM that serves a residential zone. The volume values are given as percentage of the maximum bit rate reached on the link (not the maximum capacity of the link).

Link to a business area

The traffic to a DSLAM that connects mainly business clients is characterized by a strong dependence on the type of day (Figure 67): in the working day we have the maximum link occupancy in the working hours, typically from 7-8 in the morning to 5-6 in the afternoon, with a very quick rise in the morning. In the working day we can observe also a slight traffic decrease around lunch time. In the holiday the traffic on this type of links results particularly light with respect to the traffic volume in the working day, with a maximum volume that is around 30% of the maximum in the working days.

Link to a residential area

The traffic to a DSLAM that connects mainly residential clients is not so different between holiday and working day (Figure 68). In both types of day the traffic reaches the maximum link occupancy. In the working day we can observe a slower rise in the morning, but in the afternoon the traffic levels are comparable. In the night the volume decreases to about 15-20% of the maximum.

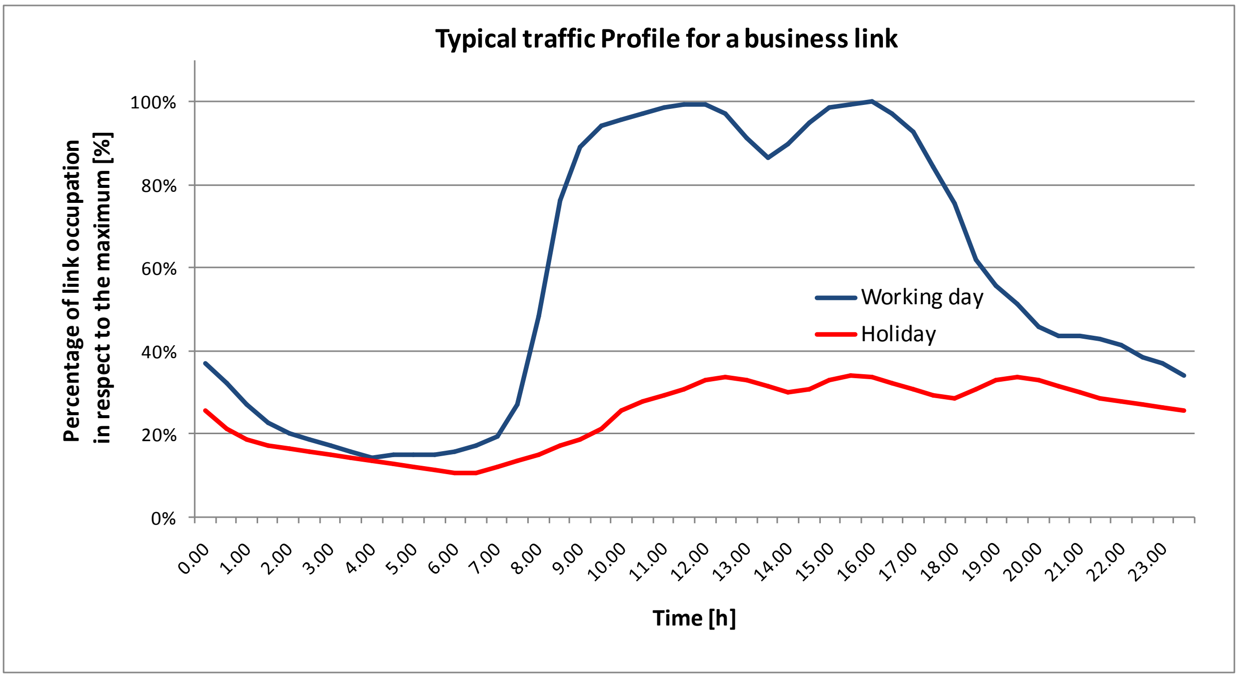

Figure 67: Typical traffic profile for a business link.

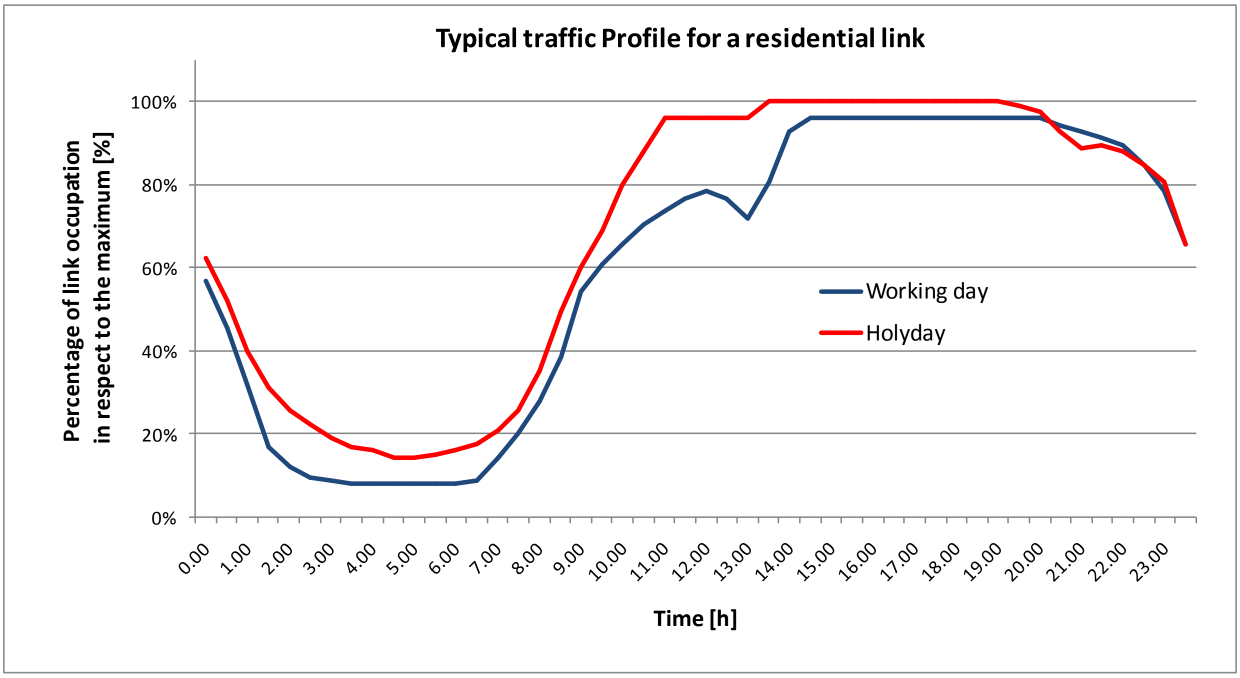

Figure 68: Typical traffic profile for a residential link.

4.2.2 Reference Internet Service Provider scenario by GRNET

The considered reference scenario by GRNET contains the number of devices in the core and access networks and includes also information regarding the data centres as well as the Grid nodes (HellasGrid – www.hellasgrid.gr).

The GRNET network

Figure 69: GRNET IP network topology.

The GRNET network is a new generation optical fibre network based on Wavelength Division Multiplexing – WDM technology at high speeds (1-10 Gbps). The core network is formed by IP routers that are interconnected with PoS 2.5 Gbps circuits over 10Gbps wavelengths that are implemented via owned DWDM equipment. Since 2008, GRNET dark fibre network is extended all over Greece, with total length of dark fibre more than 9000km and optical equipment that may support speeds up to 21x10 Gbps per link. The GRNET network topology is shown in Figure 69 and Figure 70. Figure 69 shows the GRNET IP network topology consisting of 10 IP nodes. Figure 70 shows the network topology including the established Layer 2 Ethernet links for the interconnection of GRNET clients.

The GRNET network can be divided into core and access network parts. The access network consists of dark fibre pairs between the point of presence (PoP) of GRNET in each major city in Greece and the PoP of the connected university or research institute. Around 100 clients are connected to the GRNET network. Thus, the GRNET network topology can be considered as a flat network topology without large aggregation points. Alternative backup paths are available for the majority of the network nodes while more than one alternative paths exist for the central network nodes. A more detailed view of the network is also provided in Figure 79. Regarding network congestion, this is too low to appear, except in some access links only.

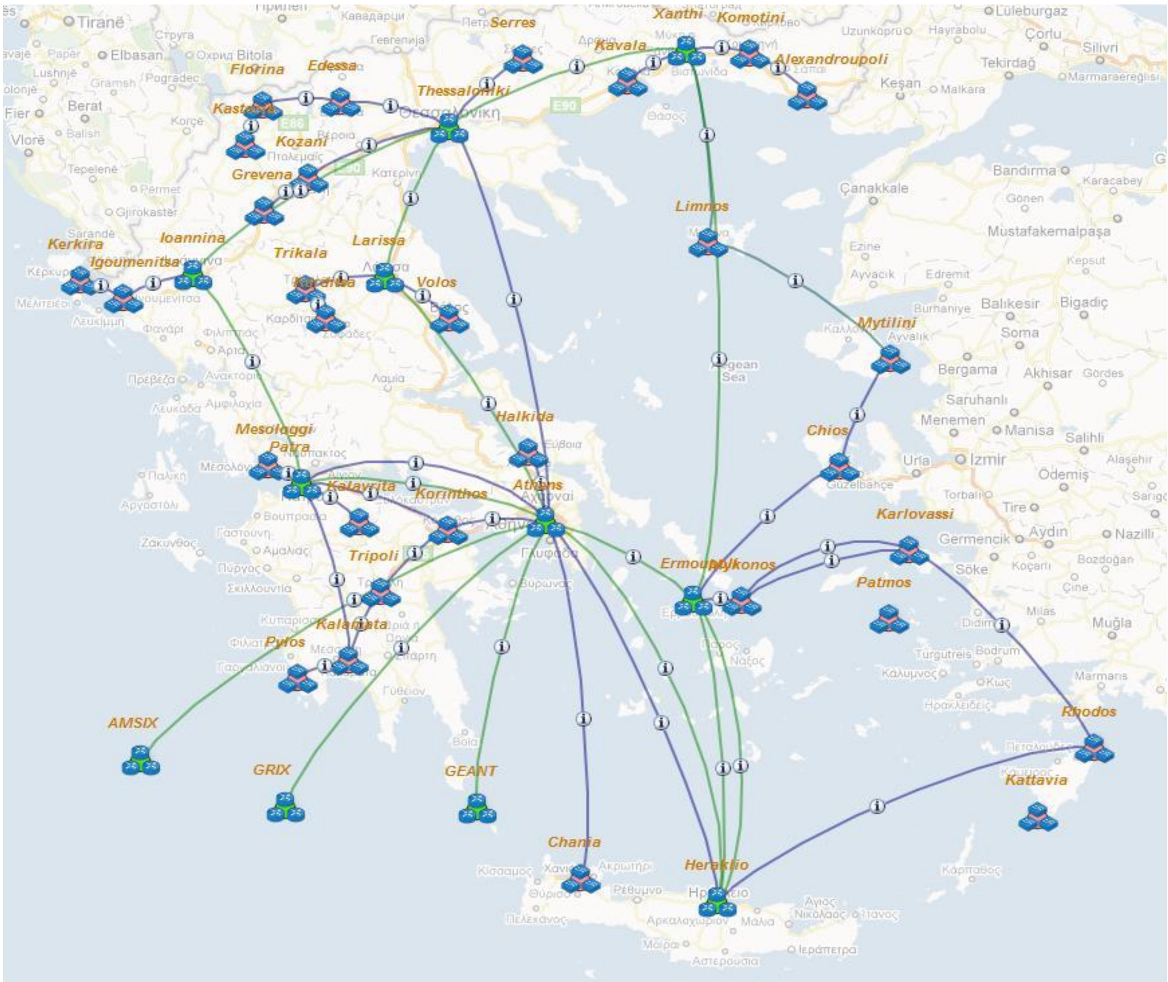

Figure 70: GRNET network topology.

Energy consumption of GRNET network equipment

GRNET's network is illustrated by IP/Ethernet and optical network equipment. The IP/Ethernet network consists of 11 routers and roughly 42 switches from different vendors. The typical energy consumption of each network device differs according to the installed cards and modules. The maximum power consumption for the access network switches and the core network routers, when fully loaded in terms of cards, modules and overall utilization, is shown in Table 17. Depending on the considered model, maximum energy consumption in the access network varies among 75 W, 190 W and 659 W, while maximum router energy consumption in the core network varies among 3000 W, 5450 W and 8000 W. Regarding the optical network, the DWDM equipment employed in the core network has a total maximum power consumption of 18,320 W. Furthermore, Table 18 defines key parameters that allow synthetically representing the average usage of network devices and links in the GRNET network.

Table 17: Maximum power consumption for the access and core network devices.

Network Part	Number of devices	Type of Equipment	Vendor/Model	Max Power Consumption
Access Network	42	Switch	Cisco Catalyst 2970, Extreme X350, Extreme X450a	190W, 75W, 659W
Core Network	9	Router	Cisco 12406, 12410, 12416	~3000W
Core Network	1	Router	Juniper MX960	5450W
Core Network	1	Router	Juniper T1600	~8000W

Table 18: Topological data (average figures) for the GRNET network. Source: Greek Research and Technology Network.

Access	redundancy degree for access devices (η_a)	13%
	redundancy degree of access links (ψ_a)	88%
Core	redundancy degree for core devices (η_c)	72%
	redundancy degree of core links (ψ_c)	71%

The power consumption of each core router has been recorded in real time for a period of 3 weeks and the results are demonstrated in Figure 71. The collected results indicate constant values in power consumptions for all the Cisco routers and small variations in the power consumption of Juniper routers. However, the observed small variations cannot be related to the amount of traffic served from a router but most probably to the starting and stopping of the device's cooling fans. A 24-hour snapshot is presented in Figure 72 providing a clearer view of the variations. Observing router EIE2's (Juniper T1600) power consumption which shows some kind of variability in conjunction to its traffic load presented later on in Figure 71, it is clear that no evident proportionality exists between energy consumption and router utilization in terms of traffic serving. The average energy consumption for each core network router is reported in Table 19 and as it may be observed router consumption ranges from 923 W to 2537 W, taking quite lower values compared to those reported before for the maximum energy consumption.

Table 19: Average power of the core network routers.

GRNET's network router	Vendor/Model	Average real time energy consumption
athens-3	Cisco 12416	2537W
eie-1	Cisco 12406	1073W
eie-2	Juniper T1600	1915W
heraklio-2	Cisco 12410	1014W
ioannina-2	Cisco 12406	939W
koletti-1	Juniper MX960	1513W
larissa-2	Cisco 12406	923W
patra-2	Cisco 12410	1154W
syros	Cisco 12406	1029W
thessaloniki-2	Cisco 12410	1428W
xanthi-2	Cisco 12406	1023W

Regarding the power consumption behaviour of GRNET's network switches, indicatively we report the measurement results from [102] on a Cisco Catalyst 2970-24TS (with 24 ports) which is a switch model highly used in our access network and which, as reported before, has a 190W maximum power consumption when fully loaded. The power consumption has been measured at the wall socket (AC) using a power meter while measurements have been taken varying the number of Ethernet links attached, the link rate of the attached interfaces (10 Mbps, 100 Mbps, 1 Gbps) and the link utilization between 0% and 100% when the attached interface's link rate was 1 Gbps; for all other rates results presented are for 100% utilization. The results are shown in Figure 73.

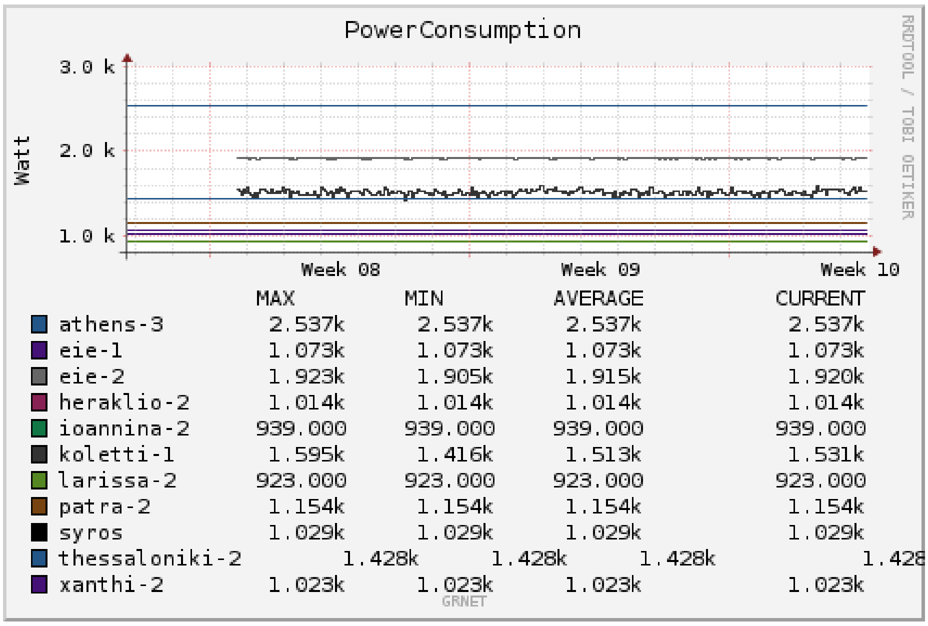

Figure 71: Power consumption in core GRNET routers over a period of 3 weeks.

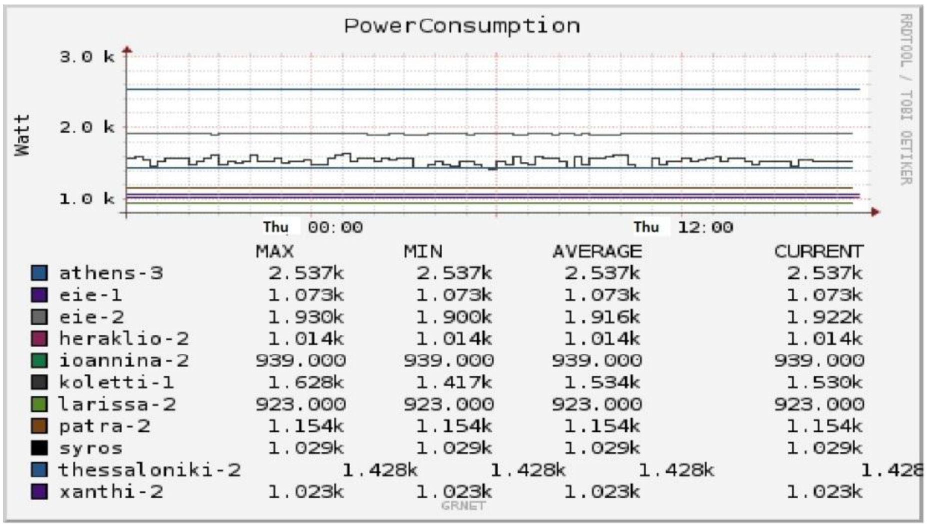

Figure 72: Power consumption in GRNET core routers for a 24-hour period.

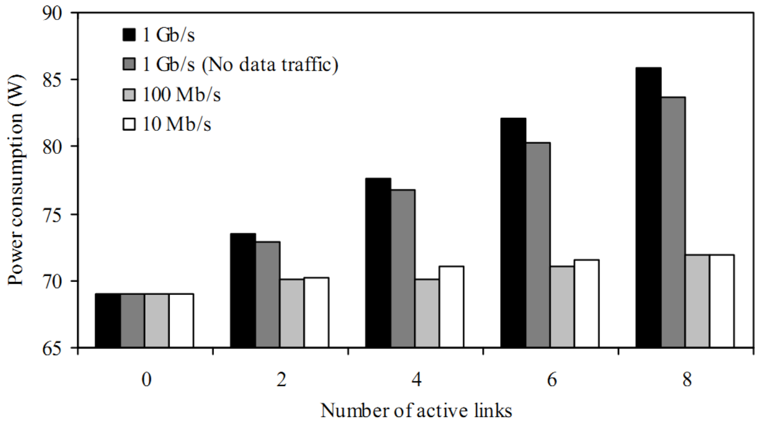

Figure 73: Power Consumption of Cisco Catalyst 2970 Switch [Source:[102]].

The observations are the following:

- The power consumption under zero utilization (no links attached) is about 68 W.
- The switch power consumption is increased by connecting a new link, even if there is no data being transmitted on this link. The increase is about 0.3 W for 10 Mbps and 100 Mbps links while it is about 1.8 W for 1 Gbps links.
- There is no significant difference in power consumption whether a port is running at 10 Mbps or 100 Mbps.
- Increasing the data rate from 100 Mbps to 1 Gbps causes switch power consumption to increase by approximately 1.5 W per link.
- The difference in power consumption is quite low when a 1 Gbps link is fully utilized compared to when it is zero utilized.

Thus, in the specific switch model significant energy savings maybe achieved only when avoiding enabling 1 Gbps link rates when not needed. A more detailed study regarding GRNET's access network switches consumptions is underway.

Energy consumption in GRNET data centres

GRNET owns two data centres where high-density computational and storage equipment is installed. The first data centre is hosted in the National Hellenic Research Foundation. This data centre hosts the GÉANT Point of Presence (PoP) in Athens as well as a HellasGrid site (Grid node). For the GÉANT PoP, there are 4 racks with servers and 14 racks with telecom equipment. These racks are fully loaded at a percentage of 60%. For the Grid node, there are 6 racks hosting servers and storage equipment. Average energy consumption for the GÉANT PoP is 63 kW; for the Grid node it is 67 kW.

The second GRNET data centre is located within the premises of the Greek Ministry of National Education and Religious Affairs in Athens. The data centre is currently equipped with 28 racks for installing servers and storage equipment. Currently, 4 racks are hosting servers and 2 racks storage equipment, but GRNET plans to load 18 extra racks in the coming months. The average energy consumption of the equipment hosted at this data centre is currently around 90 kW (see Figure 74 for total power consumption and Figure 75 for the power consumption in a rack with servers) but it is estimated to increase considerably in the upcoming period and reach 500 kW. This data centre has been designed and implemented following high standards regarding the cooling efficiency and the exact Power Usage Effectiveness (PUE) is about to be accurately determined.

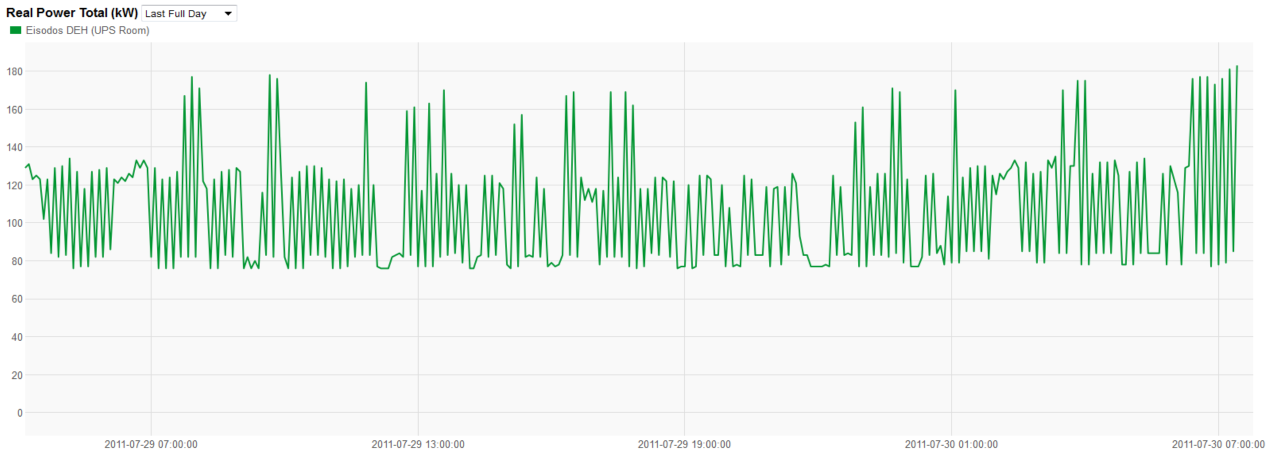

Figure 74: Power Consumption in GRNET data centre.

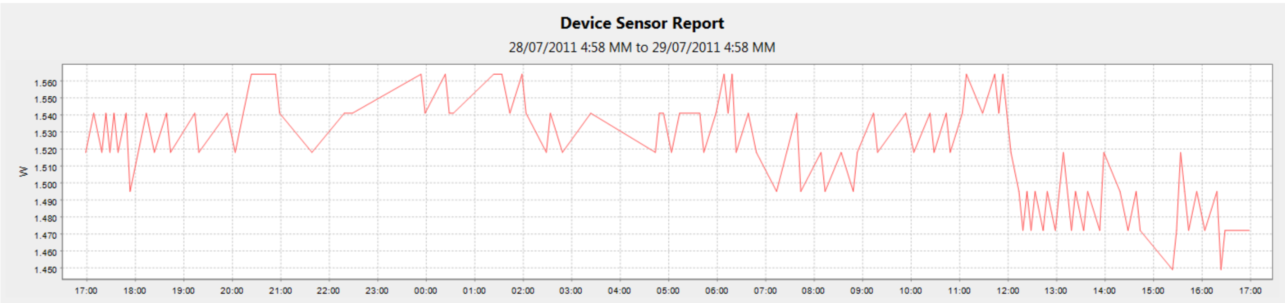

Figure 75: Power Consumption in a rack with servers in GRNET data centre.

In the slightly longer term, GRNET plans to install a green data centre outdoors in the northwest part of mainland Greece, close to a power-production hydro-electric plant facility. Water from the nearby river will be used to cool the equipment within the data centre, while for this purpose water-cooled racks will be utilized. The maximum power for the equipment hosted at this data centre is estimated to be around 400kW and the achieved PUE is expected to be among the most competitive ones. The reason for selecting water cooling is that it provides better results regarding power usage efficiency than the use of air cooling, which is the traditional data centre cooling and refrigeration technology.

Traffic profiles and volumes

A detailed profiling of the current traffic is required in order to be able to consider the potential of deploying energy efficient strategies such as rerouting traffic over energy efficient paths while shutting down others or even performing network consolidation strategies. In the following we present measurements of the link loads, while next we employ a traffic matrix estimation algorithm in order to estimate the traffic flows and come up with useful results.

Link load analysis

Link utilization as well as the monthly average and maximum throughputs are recorded. Traffic volumes per network segment are also measured. In the presented results traffic variability may be observed for selected core and access network equipment. Figure 76 shows the daily, weekly and monthly variations from the core network equipment.

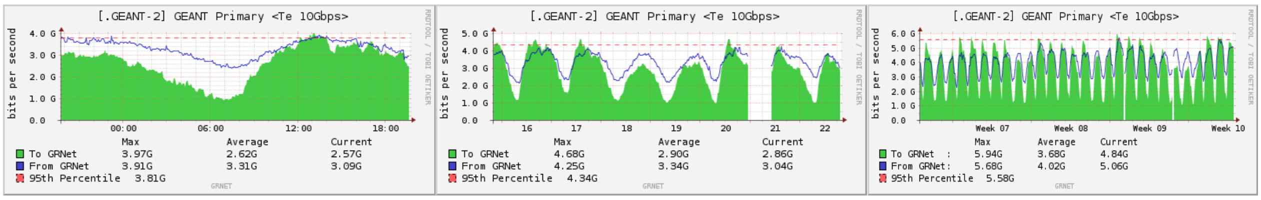

Figure 76: Traffic profile of core GRNET network router (peering with GEANT). Left: daily, Centre: weekly; Right: monthly.

Similar information is available for network routers in Figure 77, as well as access networks and GRNET power users in Figure 78.

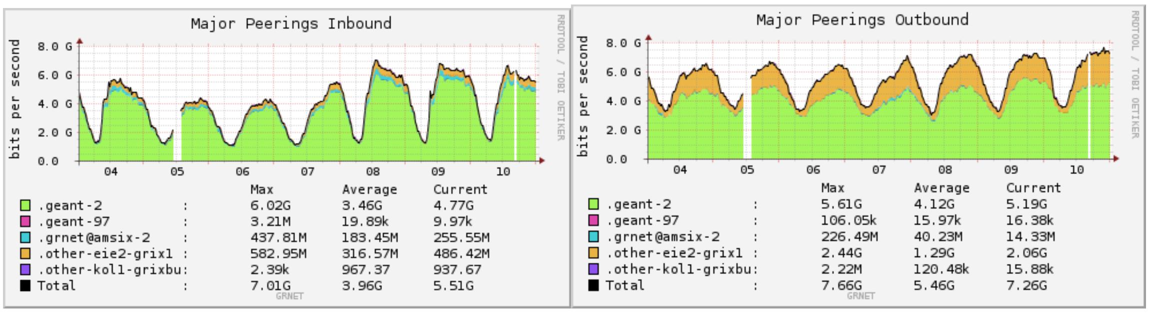

Figure 77: Daily traffic profile of core GRNET network router (Incoming and Outgoing traffic).

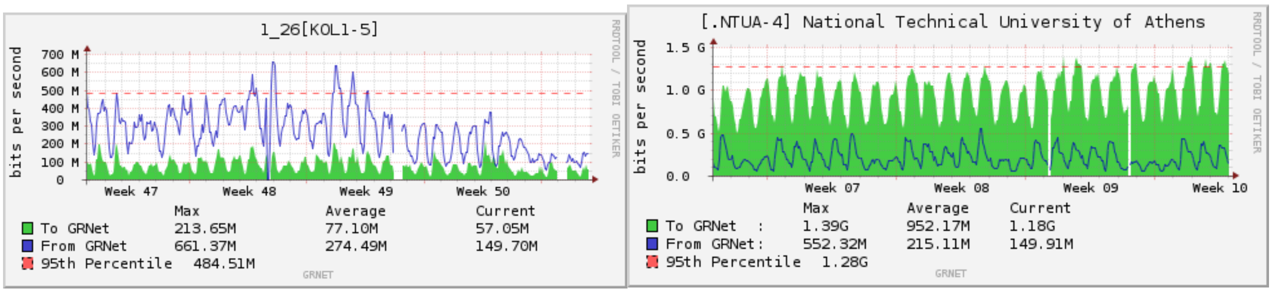

Figure 78: Daily traffic profile of access network switch (left) and a GRNET power user (client connected at 10Gbps).

Based on the collected statistics and diagrams (as shown above), it could be argued that the traffic in GRNET network nodes follows a homogeneous pattern, giving initial indications that there is room to deploy smart energy-efficient strategies and thus improve the energy-efficiency of the network.

Flow level analysis

In the absence of a unified method for data flow statistics collection in GRNET network, a traffic matrix estimation methodology is applied having as input the distinct link loads so as to be able to determine network flows between all source-destination pairs. This estimation procedure is expected to be assisted in the future by partial flow-level measurements in order to produce more accurate estimates. In the following first we present the traffic matrix estimation approach while then we provide answer to a set of statistical questions regarding the origins and destinations of traffic as well as regarding transit traffic.

Traffic matrix estimation

A traffic matrix (TM) provides, for every ingress point i into the network and every egress point j out of the network, the volume of traffic X_{ij} from i to j over a given time interval. Traffic matrices are critical inputs to network design, capacity planning, reliability analysis, traffic engineering and business planning.

The TM may be computed in a direct manner by collecting flow-level measurements at ingress points of the network. However, this turns into a great burden for network providers since it requires collecting a huge amount of traces per day while it demands additional features on all routers and the whole collection procedure often reduces a router's forwarding performance. These drawbacks along with other practical and compatibility problems in infrastructures with diverse equipment, make network wide flow-level measurements difficult if not impossible. This fact has triggered research in traffic matrix estimation using easily obtainable information. Information utilized from proposed TM estimation methods in order to produce their estimates consists of:

- link loads (counts) most commonly obtained in every network with SNMP,
- routing information (weights on links) that permits producing the routing matrix of the network,
- additional topological information regarding the role of each node (e.g. peering node, access node) and
- assumptions on the distribution of demands.

Let there be n network nodes and r directional links, then there will be $k = n \cdot (n - 1)$ source – destinations (SD) pairs. Let X be the traffic matrix with dimensions $k \times 1$ where x_p (element at the p^{th}row of the matrix) is the data transmitted by the SD pair $p \in [1..k]$. Let Y be the matrix of link loads (counts) with dimensions $r \times 1$ where y_l (element at the l^{th} row of the matrix) is the load of link $l \in [1..r]$. Let A be the routing matrix with dimensions $r \times k$ where $a_{l,p}$ (element at the l^{th} row, p^{th}column of the matrix) is equal to 1 if link l belongs to the path associated to SD pair p, otherwise is equal to 0. Then the following equation holds:

$$Y = A \cdot X \quad (1)$$

In TM estimation equations Y and A are known while X is the unknown. This matrix equation is under constrained (ill-posed) since $r < k$ allowing many solutions. The TM estimation methods base their operation in figuring out which TM is the most likely one. Popular methods include linear programming, which introduces new objective functions and/or constraints and solves an optimization problem to decide an estimate of the TM [103], and methods that make assumptions on the statistical distribution of the TM in order to develop TM estimation techniques [103]. A noticeable development in the area is the introduction of the gravity models and related techniques [105][106].

Let the total inbound flow N_s^{in} from a source node s be the sum of the loads over all the inbound edge links from s and the total outbound flow N_d^{out} to a destination node d be the sum of the loads over all the outbound edge links to d. It has been observed [104] that for each fixed source node s, the distribution of the inbound traffic N_s^{in} from s to different destinations d is approximately proportional to the total outbound loads N_d^{out} these destinations receive. Formally, this is called the gravity model and can be written as:

$$\tilde{x}_j = \frac{N_{s_j}^{in} \cdot N_{d_j}^{out}}{N} \qquad (2)$$

where $N = \sum_s N_s^{in} = \sum_d N_d^{out}$, s_j,d_j are respectively the source and destination nodes for the j^{th} SD pair and $\tilde{x}_j$ is the corresponding component of the simple gravity solution $\tilde{X}$.

$\tilde{X}$ matrix is an approximation of the traffic matrix but clearly does not satisfy the constraints for the link loads (eq. 1). A simple tomogravity model performs a two-step modelling:

- At the first step, it obtains an initial solution, by applying the Gravity Model.
- At the second step, it performs a tomographic estimation, according to which the initial solution is refined by applying quadratic programming to minimize the distance to the initial solution, subject to tomographic constraints (link counts).

The simple tomogravity solution is the result of the following optimization problem:

$$\text{argmin}_X\{\|X - \tilde{X}\|: A \cdot X = Y\} \quad (3)$$

where $\|\cdot\|$ denotes the $l^2 - norm$ of a vector[6].

More general tomogravity solutions have been introduced, where the edge nodes are further classified as “access” or “peering” [104]. According to [104] tomogravity models have been found to perform quite well, with an average error in estimation of about 12%. Furthermore, if there are known TM rows from direct flow measurements these may highly assist the model to produce a more accurate solution and reduce the estimation error.

Traffic matrix estimation procedure

We implemented the simple tomogravity method in order to produce the GRNET network TM. The full network topology consisting of Layer 2 and Layer 3 parts is shown in Figure 79. In order to capture the essential information regarding traffic flows, while reducing the size of the problem we consider the topology shown in Figure 80 which is the product of a procedure that performs:

- geographical aggregation,
- removal of the alternative backup links while
- preserving backbone links and
- considering power users.

[6] For a real vector $X = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$, $\|X\| = \sqrt{{x_1}^2 + {x_1}^2 + \cdots + {x_n}^2}$

Then, for the abstracted topology shown in Figure 80 the routing matrix A is computed performing Dijkstra's shortest path algorithm with the weight actually used in our network for each link equal to $\frac{link\ capacity\ (Gbps)}{100\ (Gbps)}$.

The average monthly loads of the links of the network have been gathered[7] and are used to form the Y matrix.

Then, we consider solving, using OPL CPLEX[8], a more constrained version of the optimization problem presented in eq. 3, in order to further restrict the solution space and get a good solution in shorter time. The details are provided in the following.

Let $\tilde{X}$ be the simple gravity solution computed for the GRNET network. Let X' be a matrix with dimensions kx1 where x'_p (element at the p^{th}row of the matrix) is equal to -1, if the corresponding row of the traffic matrix is not known from flow measurements or known policy and equal to the actual flow volume, if it is known. We consider that every network node has information to exchange with every other node of the network unless this is not allowed by the policy followed. If some nodes do not communicate with each other due to the network policy or if some flows have been actually measured, this is expressed by matrix X'. Furthermore, we have taken into account that the volume of a flow from an origin to a destination will not be higher than the minimum load of a link in the path from the origin to the destination. The optimization problem solved follows.

$$\begin{aligned}
&\textit{minimize } \|X - \tilde{X}\| \\
&\textit{subject to} \\
&\quad A \cdot X = Y \\
&\quad x_p = x'_p \ \forall\, p: x'_p \neq -1 && \text{// known traffic matrix rows} \\
&\quad x_p > 0 \ \forall\, p: x'_p = -1 && \text{// lower bound} \\
&\quad x_p \leq \min(a_{l,p} \cdot y_l \ \forall\, a_{l,p} \neq 0) && \text{// upper bound}
\end{aligned}$$

The solution to the above problem giving an estimated TM for the GRNET network has been estimated, taking into account that GÉANT, GRIX[9] and AMSIX[10] according to the followed policy do not exchange any volumes with each other.

Outcomes

The estimated TM for the GRNET network assists us in extracting a series of quantitative results and consequent qualitative results regarding the network traffic. These results are a useful input when either reconsidering network design or considering traffic engineering approaches especially when this is done from the energy aware perspective.
Given the estimated TM, the following questions are considered:

---

[7] http://mon.grnet.gr/

[8] http://www-01.ibm.com/software/integration/optimization/cplex-optimization-studio/

[9] GRIX is the Greek internet exchange node providing interconnectivity between Greek ISPs, http://www.gr-ix.gr/
[10] AMSIX is the Amsterdam Internet Exchange, http://www.ams-ix.net/

1. ***What is the percentage of the incoming traffic to a node that terminates at the node?***

Let $x_{i\to k}$ be the traffic flow volume from a network node i to node k and $y_{j\to k}$ be the downstream traffic load of a link between node i and node k. For a node k the requested ratio is expressed as:

$$r_1 = \frac{\sum_i x_{i\to k}}{\sum_j y_{j\to k}}$$

2. ***What is the percentage of the incoming traffic to a node that traverses (passes through) the node?***

For a node k the requested ratio is expressed as:

$$r_2 = \frac{\sum_j y_{j\to k} - \sum_i x_{i\to k}}{\sum_j y_{j\to k}}$$

3. ***What is the percentage of outgoing traffic from a node that originates from the node?***

For a node k the requested ratio is expressed as:

$$r_3 = \frac{\sum_i x_{k\to i}}{\sum_j y_{k\to j}}$$

4. ***What is the percentage of outgoing traffic from a node that traverses the node?***

For a node k the requested ratio is expressed as:

$$r_4 = \frac{\sum_j y_{k\to j} - \sum_i x_{k\to i}}{\sum_j y_{k\to j}}$$

The following stands:

$$r_2 \cdot \sum_j y_{j\to k} = r_4 \cdot \sum_j y_{k\to j}$$

5. ***What is the percentage of the total traffic served by the node that terminates at, or originates from the node?***

For a node k the requested ratio is expressed as:

$$r_5 = \frac{\sum_i x_{i\to k} + \sum_i x_{k\to i}}{\sum_j y_{j\to k} + \sum_i x_{k\to i}}$$

where $\sum_j y_{j\to k} + \sum_i x_{k\to i} = \sum_j y_{k\to j} + \sum_i x_{i\to k}$ is the total traffic served by the node consisting of the traffic originated by the node, the traffic terminated to the node and the traffic that traverses the node.

6. What is the percentage of the total traffic served by the node that traverses the node?

For a node k the requested ratio is expressed as:

$$r_6 = \frac{\sum_j y_{k\to j} - \sum_i x_{k\to i}}{\sum_j y_{j\to k} + \sum_i x_{k\to i}}$$

7. What is the percentage of the total network traffic that is originated by a node?

For a node k the requested ratio is expressed as:

$$r_7 = \frac{\sum_i x_{k\to i}}{\sum_i \sum_j x_{i\to j}}$$

where $\sum_i \sum_j x_{i\to j} = \sum_i \sum_j x_{j\to i}$ is the total network traffic as the sum of all network flows.

8. What is the percentage of the total network traffic that is terminated at a node?

For a node k the requested ratio is expressed as:

$$r_8 = \frac{\sum_i x_{i\to k}}{\sum_i \sum_j x_{j\to i}}$$

9. What is the percentage of the total network traffic that traverses this node?

For a node k the requested ratio is expressed as:

$$r_9 = \frac{\sum_j y_{j\to k} - \sum_i x_{i\to k}}{\sum_i \sum_j x_{j\to i}}$$

The aforementioned ratios are computed and the results lead to the following selected observations:

- *Regarding the origins of the traffic*
    - ~32% of the total network traffic is originated in Athens (ATH and power users NTUA, UOA) and ~33% of the total network traffic is incoming traffic from GÉANT. Significant percentages of traffic are also originated in campuses or cities LANs-MANs (4.09% in HERSW1, 5,.2% in PATSW, 1.95% in NTUA, 1.9% in XANTHISD). Traffic originated from GRIX is 2.83% of the total network traffic.
- *Regarding the destinations of the traffic*
    - ~36% of the total network traffic is terminated in Athens, ~35% of the total network traffic is outgoing traffic to GÉANT and ~10.5% of the total network traffic is outgoing traffic to GRIX. Traffic terminated in NTUA is 8.13% of the total network traffic.
- *Regarding transit traffic*
    - The transit traffic in Athens is ~40% of the total network traffic. The corresponding percentages for the rest core IP router nodes are below 4% (3.59% for LRS, 3.14% for PAT, 2.87% for THES, 2.77% for SYR, 2.35% for XNT, 0.59% for IOA and 0.01% for HER). The remaining nodes in the network do not carry transit traffic.

- In Athens (ATH and power users NTUA, UOA), 54% of the total served traffic by the node is forwarded to other nodes (transit traffic). The other core IP routers (except HER) carry transit traffic in percentage that varies from 20% to 80% of each router's total served traffic; in Patra (PAT) (56%), Ioannina (IOA) (22%), Xanthi (XNT) (80%), Thessaloniki (THES) (63%), Larisa (LRS) (67%) and Syros (SYR) (72%). The router in Patra forwards transit traffic passing through the nodes ATH←→IOA, the router in Ioannina transit traffic passing through the nodes IOA←→THES, the router in Thessaloniki transit traffic passing through the nodes LRS←→XNT, LRS←→IOA and IOA←→XNT, the router in Xanthi transit traffic passing through the nodes SYR←→THES, the router in Syros transit traffic passing through the nodes ATH←→HER and XNT←→HER, and the router in Larisa transit traffic passing through the nodes ATH←→THES. The core router in Heraklion (HER) does not carry transit traffic.

- *Regarding geographical distribution of the traffic*
    - If the Athens node is abstracted so as to include capacity demanding local customers (NTUA, UOA) as well as third party networks that interconnect with the GRNET network at the Athens area (GEANT, GRIX, AMSIX) then the analysis show that 70.18% of the overall network traffic is produced at the extended Athens area. From this amount, 80.88% of the traffic is internally consumed and the rest 19,12% of the traffic is directed outside the Athens area. Moreover the amount of traffic that is directed to the extended Athens node from customers outside Athens accounts for 23.5% of the total traffic. Hence 93.23% (70,18+23,05) of the total network traffic is either produced at the extended Athens node or directed to it from the rest of the network.

Clearly, there is a high percentage of aggregation of traffic in Athens while in comparison the remaining routers may be considered as underutilized. Furthermore, traffic from and to GÉANT and GRIX is almost 81% of the total traffic. The obtained outcomes are considered as a valuable input that will drive our energy aware traffic engineering study.

Concluding remarks and energy efficiency perspective

Results have not revealed any kind of energy proportionality to the traffic served by a network device. However, it seems that there is space for energy gains by altering link speeds or by even rerouting traffic and shutting down network equipment according to usage patterns. Furthermore, it seems that network consolidation strategies may also be highly applicable. However, employing such approaches requires several other considerations. Towards our next step to consider energy aware traffic engineering strategies, power consumption measurements along with detailed link utilization measurements and the estimation of the traffic matrix assisted by partial actual flow-level measurements are considered as valuable tools.

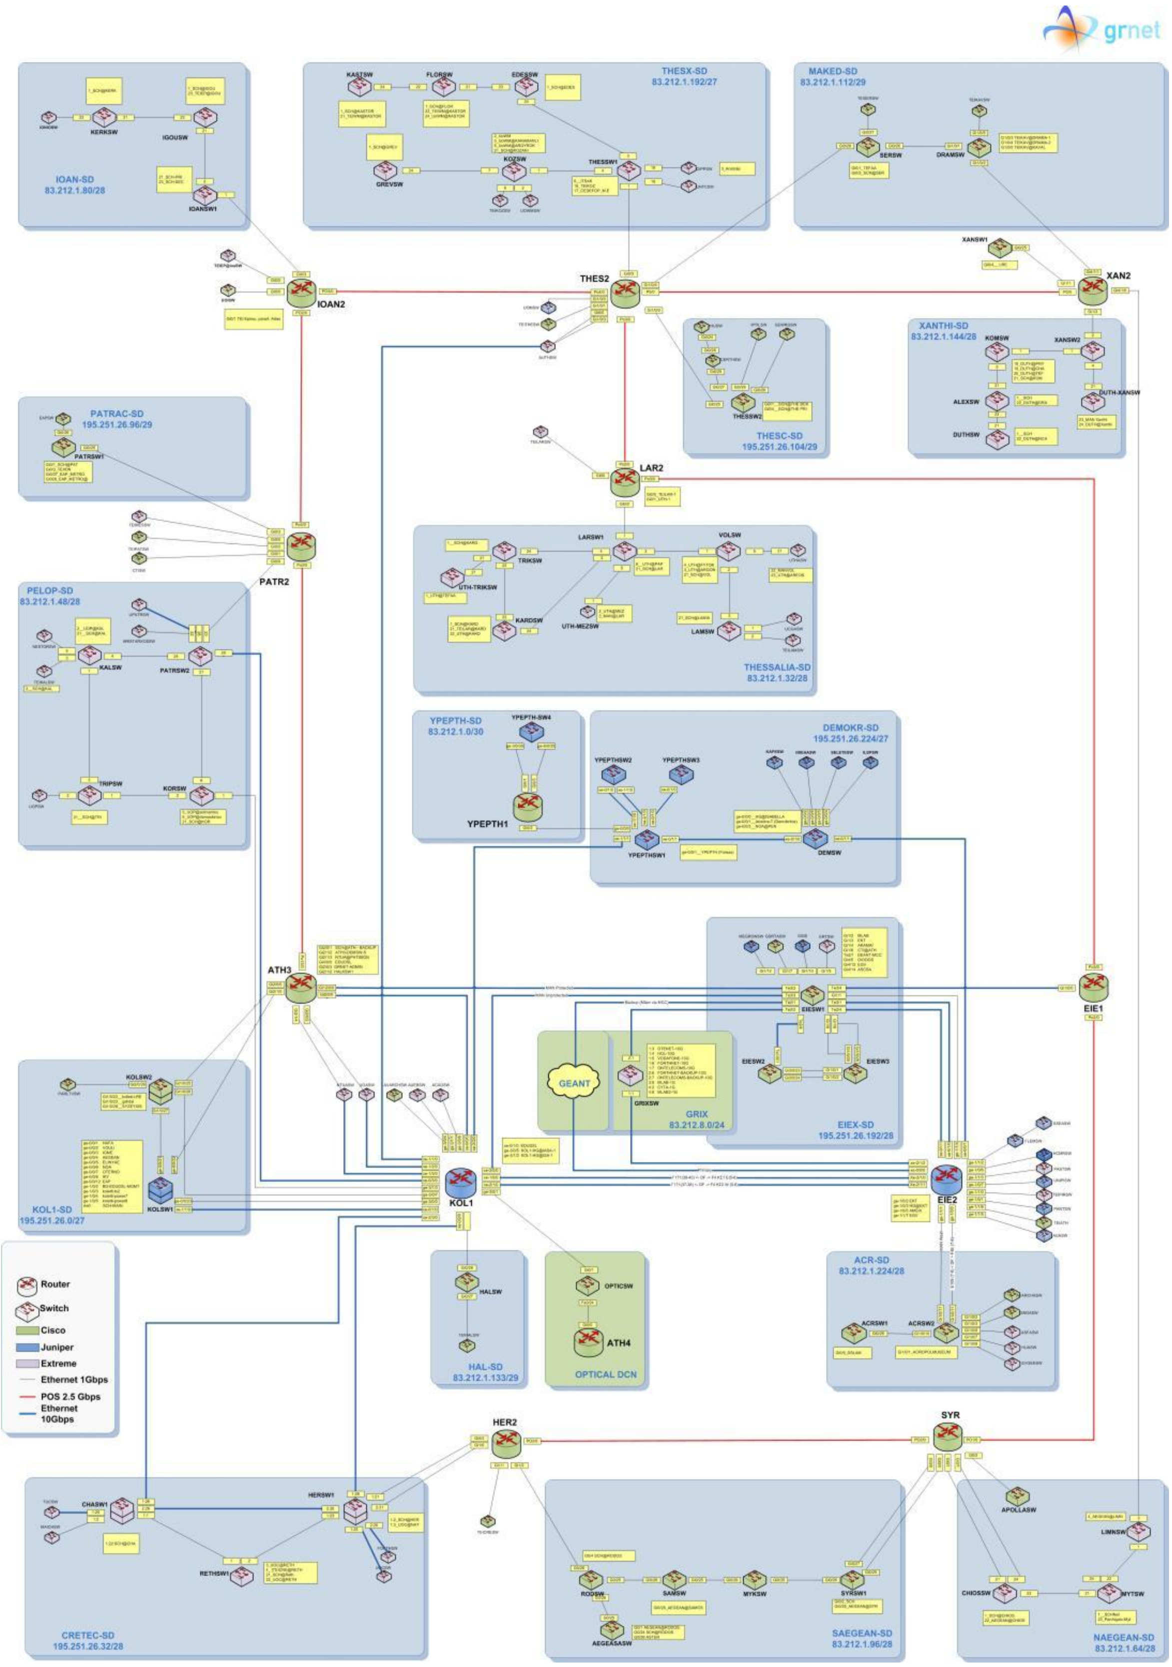

Figure 79: GRNET L2 and L3 network topology.

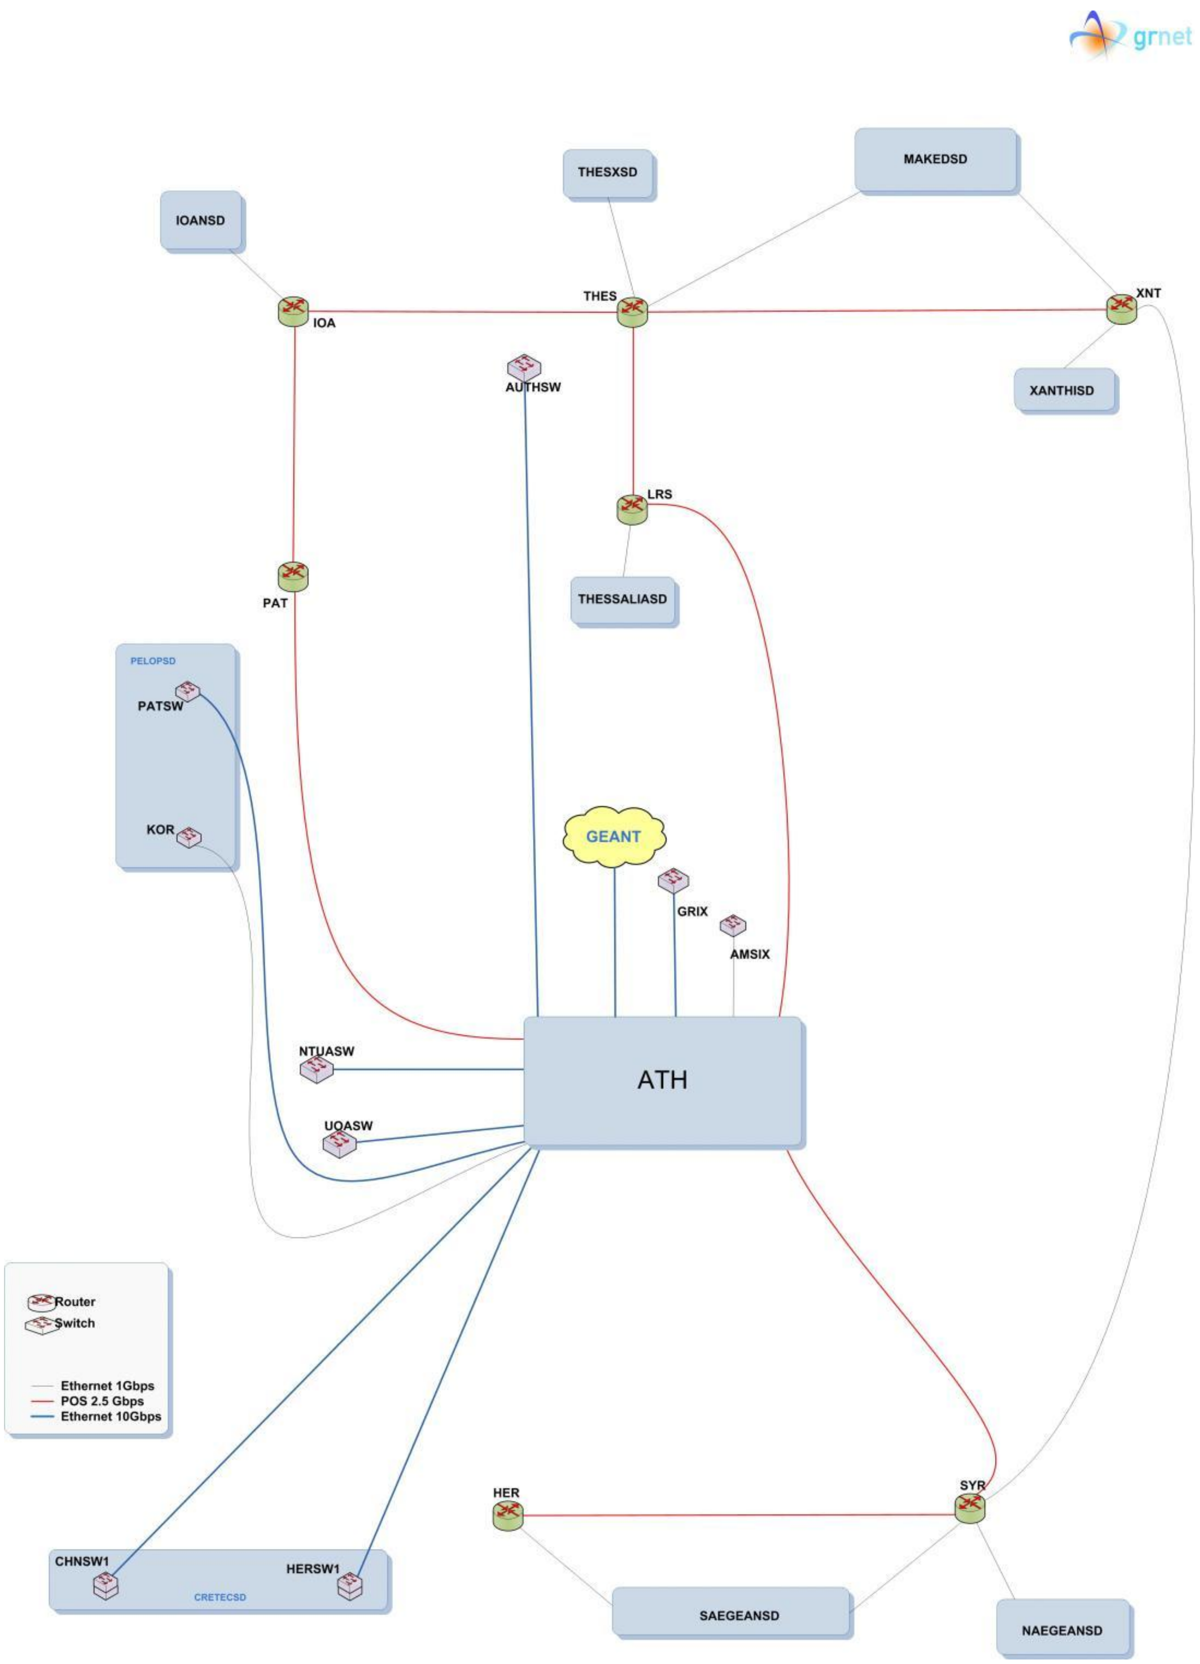

Figure 80: GRNET network topology abstraction.

4.2.3 Reference Internet Service Provider scenario by NASK

The NASK network

The networking infrastructure owned and administrated by NASK is comprised by three networks of different kinds:

- Metropolitan area network in Warsaw, connecting mainly academia and public institutions,
- Country-wide area network, connecting wired and wireless access points in major Polish cities as well as locations in rural regions supported by governmental programmes,
- Wireless 3.5-3.7 GHz base stations in major Polish cities and in the rural regions.

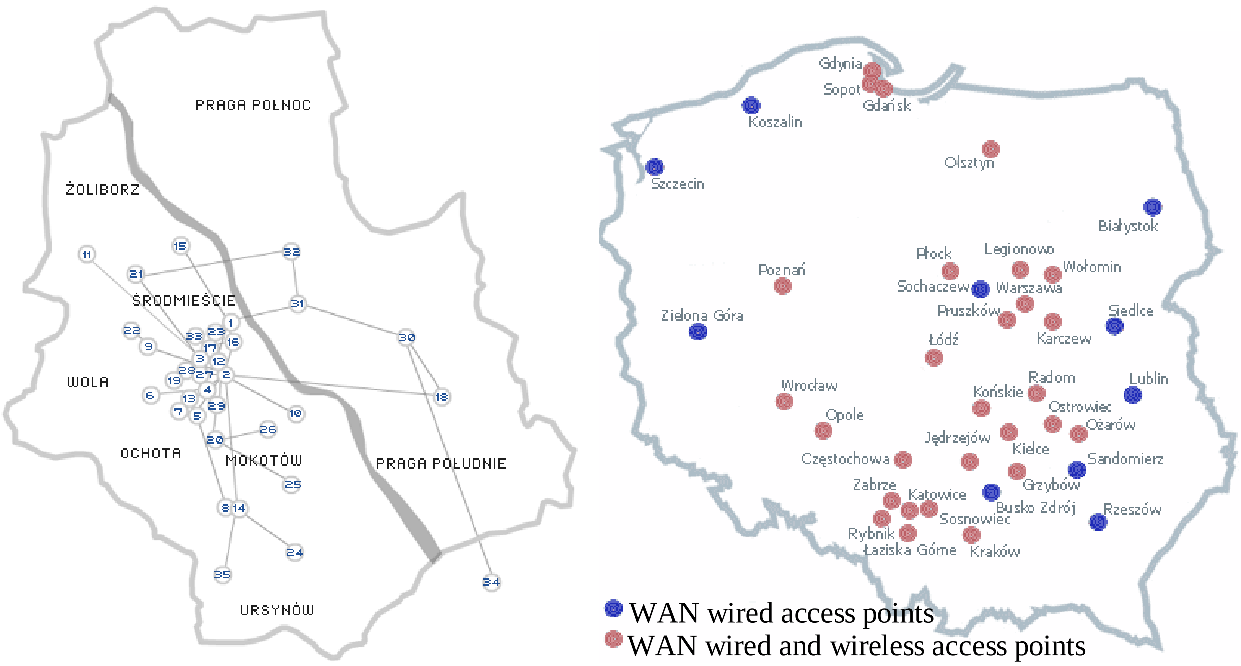

Figure 81: Topology of metropolitan (left), WAN and wireless (right) NASK networks.

The reference scenario concerns the metropolitan network, although many of the observations given below are valid also for WAN and wireless networks. The Warsaw metropolitan network has been chosen because of its adequate complexity, substantial link capacities and networking equipment capabilities. Built with gigabit switches, it makes possible to easily profit from green technologies both in data and control planes. It would become the first field where any new power saving equipment could be installed.

The profitability of any new technology to be deployed in an ISP's network is, as in any other case, a simple balance of the technology costs and benefits. In the case of the technology to be developed in the ECONET project, the final costs are still unknown. Even with green network cards prototypes developed, the final cost will depend on the production volume of their marketable counterparts. NASK, being a small operator, can benefit best only if the green networking technology is mature and relatively cheap, unlike for big telcos, capable of generating huge demand for the cards and negotiating good prices.

Whatever the costs will be, it is possible now to estimate benefits from the deployment of ECONET technology. If the power (and cost) reduction is to be perceptible, the power costs themselves must first be measurable. Simple as it may seem, it is often not the case in the NASK metropolitan network. A part of its points of presence (POPs) operates in locations where no energy accounting takes place. It is even hard to estimate the part of the contracted rental lump sum that shall be attributed to energy costs. Therefore, ECONET technology deployment should take place only there where energy consumed is properly metered, unless one cares only for global carbon footprint.

Another deployment criterion is the possibility of ECONET technology installation on both ends of a network link. This is not a strict requirement – otherwise going green at one node would imply doing it in the whole network; however, having compatible equipment at both link ends gives us more saving from low-level traffic control procedures. Alternatively, connecting green and non-green devices we can mainly gain from putting the link down altogether, when it is not really needed.

Finally, the new technology should be deployed first on the links where the potential for energy reduction is high. A good approximate of the potential is unused link bandwidth, expressed in absolute terms. The unused bandwidth is calculated as the difference of link capacity and the link utilization, averaged over a long horizon, i.e. in the order of weeks. Such measure is proposed on the assumption that a green device adapts to any traffic change: in practice there is the interplay of traffic and local control strategy dynamics, and the accurate measure of potential savings should be more complex.

NASK network equipment and typical energy consumption

The major part of NASK Warsaw metropolitan network is composed of Juniper and Cisco Ethernet switches, operating in 1 GB copper and fibre optic links. In residual parts of the network ATM technology is still in use. Below we give the most frequently used equipment models in NASK network and their declared power requirements for the respective configurations:

- Juniper M7i switch – 378 VA
- Juniper M10i switch – 576 VA
- Juniper M120i switch – 2150 VA
- Cisco Catalyst 6500 switch – 1300 VA
- Cisco Catalyst 8540 switch – 1760 VA

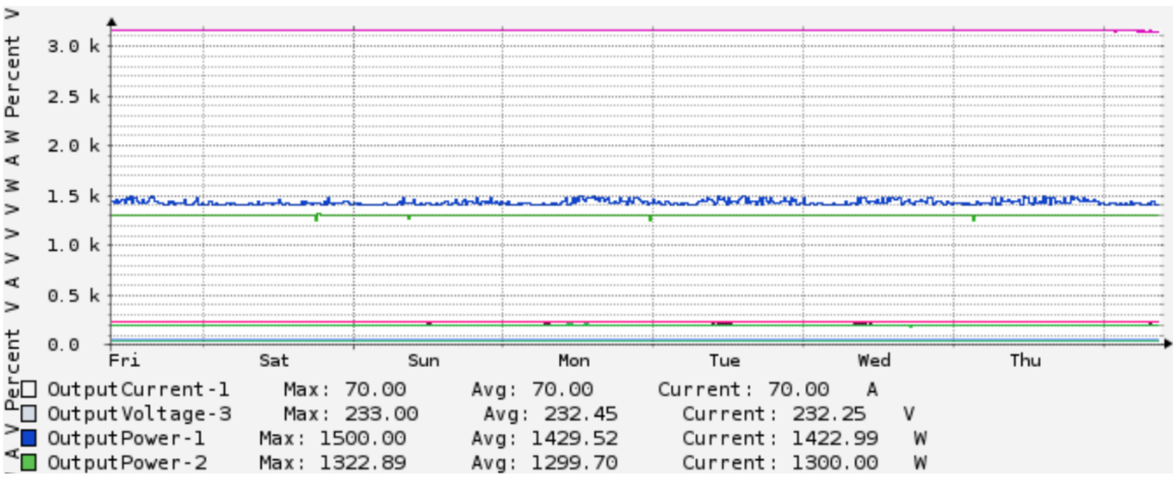

Figure 82: Power consumption in NASK network devices.

Typical graphs of power load reported by a UPS are presented in Figure 82. By observing the "OutputPower-1" and "2" graphs, one can learn that some devices exhibit slight daily variation in power consumption, weakly but evidently correlated with traffic volume, while others do not. The conclusion can be that there are already standard low-level Ethernet green features active in the network but the degree of savings is almost negligible.

NASK traffic profiles and volumes

The first link to be described goes to the NASK hosting and collocation centre where energy consumption is, naturally, properly measured. Traffic profiles for various time horizons are presented in Figure 83. The link selected is generally slightly loaded (w.r.t. link capacity of 1 Gbps), although the long-term trend is going up. The statistics show huge load variations, from almost zero to some 100 Mbps (which is probably the bottleneck in some access network of NASK academic customers). However, relative traffic load is so light that it might be argued if ECONET power scaling features would be able to handle such subtle changes, i.e. by a small fraction of link capacity. On the other hand, small traffic load prompts to apply ECONET distributed control strategies, and to route this traffic around using links that are already considerably loaded.

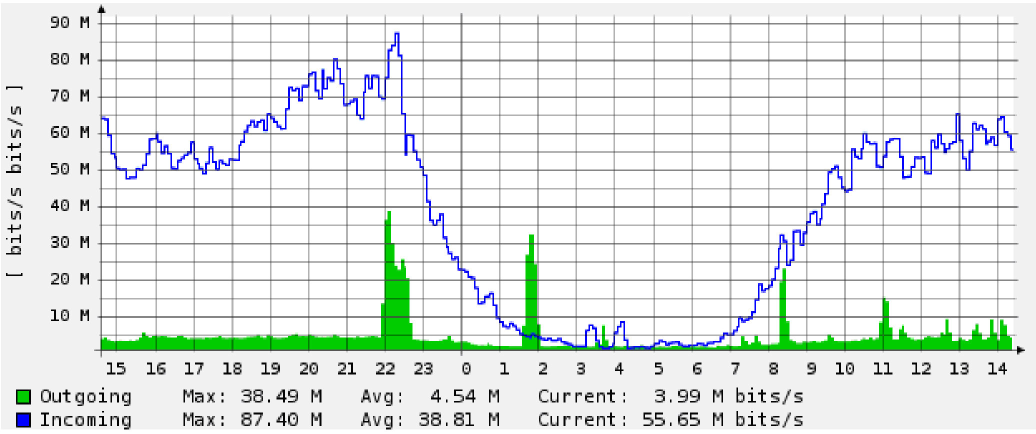

a)

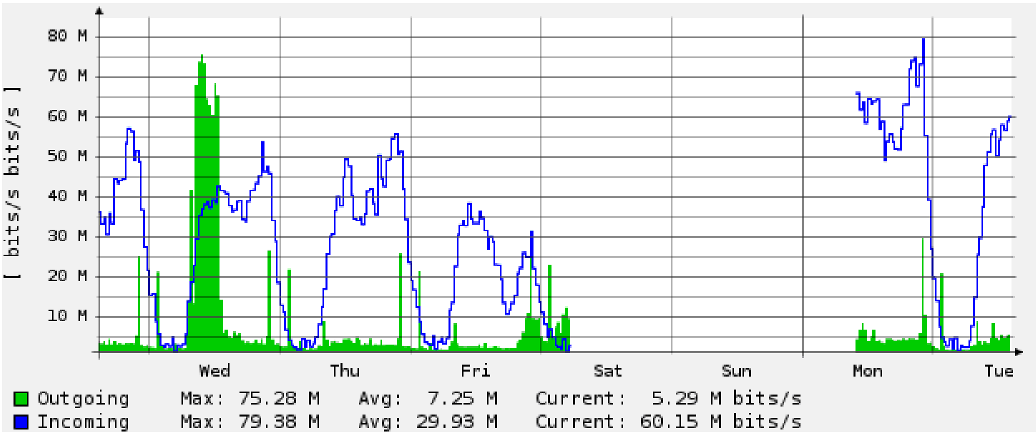

b)

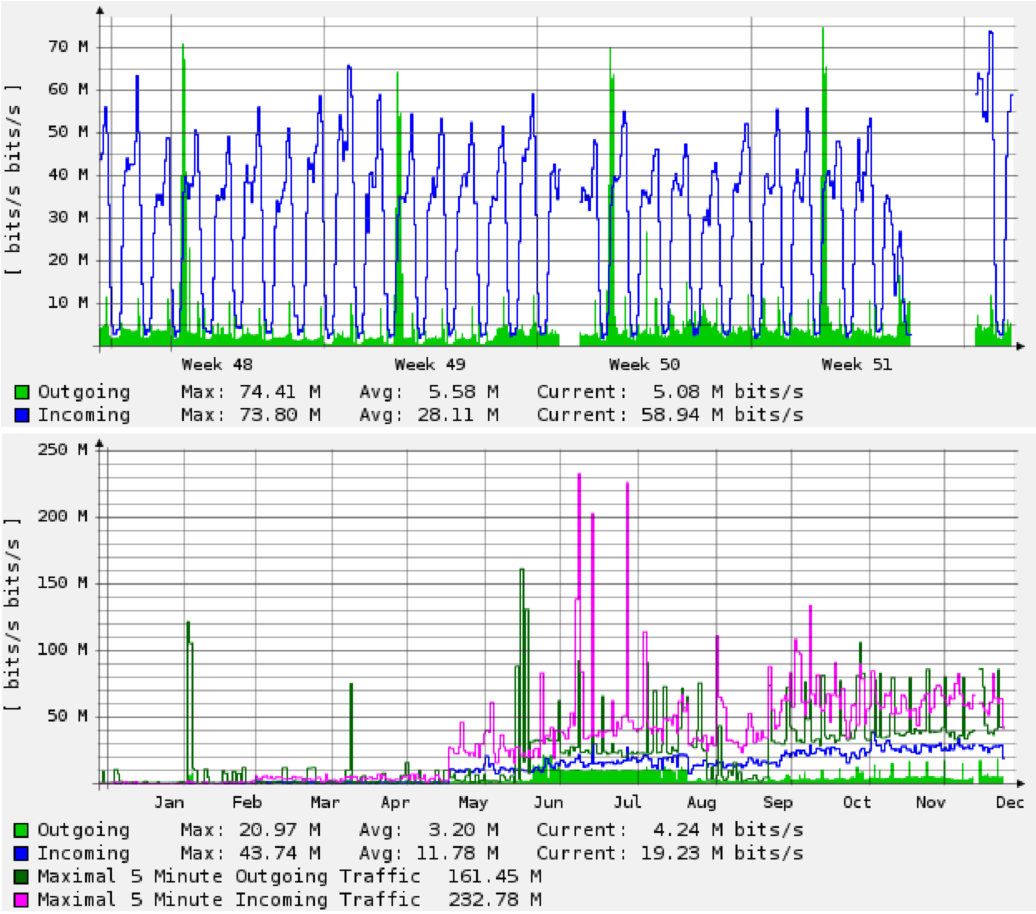

c)
d)

Figure 83: Daily (a), weekly (b), monthly (c) and yearly (d) traffic profiles.

Another link is the interconnection to a big Polish telecommunication operator. Therefore, the link goes outside the NASK network, so proper energy measuring at the far end is irrelevant. Much more important is the ability of both operators to deploy green equipment that enables efficient power scaling. Observation of Figure 84 shows a big variability of the traffic, like for the link described previously. Also, the link utilization is much higher than previously, and the peaks can reach up to 90 % of the link capacity. But generally the traffic is much smoother, probably because statistical multiplexing takes place. One should also observe occasional gaps in statistical data, including the weekend of Dec. 25/26. Such white spots could be also observed in previous series of graphs in Figure 84.

Summarizing, neither lightly loaded nor interconnection links should prevent us from green technologies deployment. However, the current business attitude towards energy saving is far from the desired state: nowadays it is too often assumed by NASK contractors that energy consumption is constant and, if so, it needs not to be measured in any way. If it were measured, it would affect the algorithm splitting common costs across services.

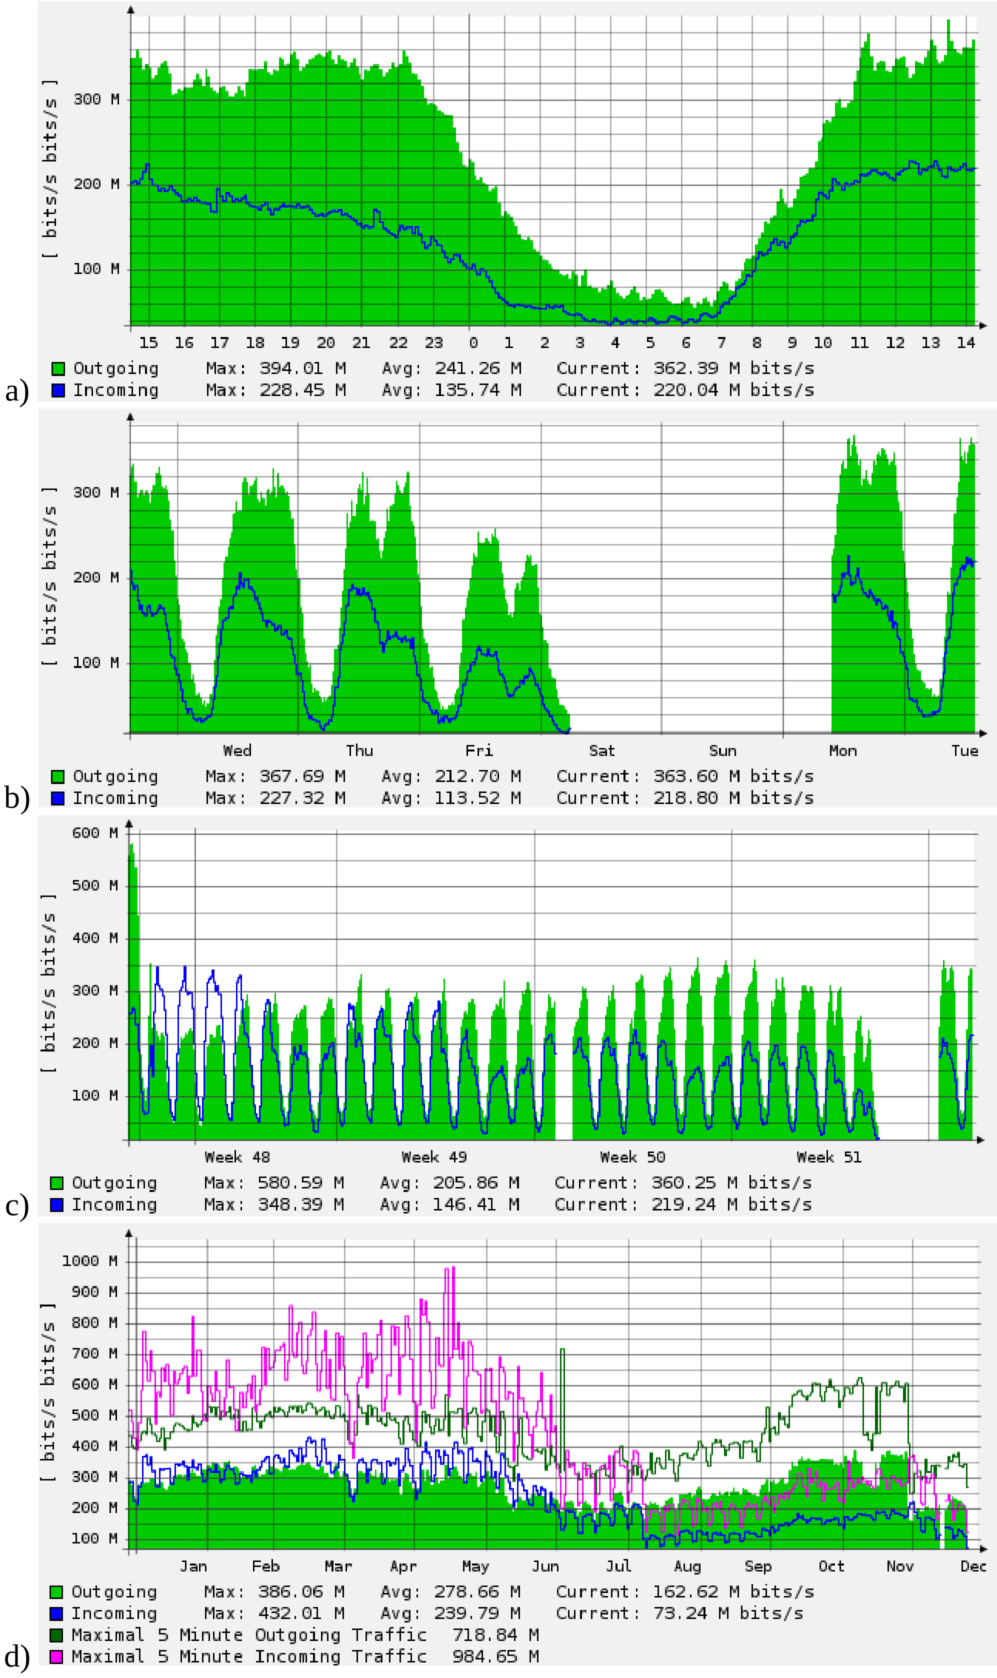

Figure 84: Daily (a), weekly (b), monthly (c) and yearly (d) traffic profiles.

Energy costs affect, as any other costs incurred, the final cost of all services offered by NASK. A complex cost assignment strategy is currently applied to calculate costs of services. One particular innovation of the cost assignment strategy is how common costs are split across services – given the fact that those services often rely on each other.

First, services are qualified to four major layers: transmission, ATM, Ethernet and IP. Likewise, network devices located in every PoP can be qualified to those service layers. For a given POP, let us denote total depreciation of devices in layer $l \in$ {transmission, ATM, Ethernet, IP} by d_l. The ratio $d_l / \sum_i d_i$ of depreciation in layer l to the total depreciation of all network devices (i.e. those that can be assigned to any service layer) determines how the cost non-applicable to any layer will be split across service layers. Such non-applicable costs include rent, maintenance (general, UPS, cabinets, air-condition etc.), depreciation (UPS, cabinets, air-condition) and electricity costs.

In the current setting, electricity costs affect costs in all service layers in proportions depending on the equipment acquisition cost. Any cost reduction from ECONET technology application will therefore be properly allocated geographically, as the common costs splitting is performed per PoP – but it will affect all service layers regardless of the service where actually savings have been done. This will lead to cross-financing of services: those provided with energy-efficient technologies will support energy-inefficient ones. Thus, any investments in green technologies will affect economy of network operation only by their acquisition costs.

Such approach must be subject to change. Otherwise, saving allowed by the application of some energy efficient technology, bought at decent prices, will not significantly decrease costs of networking services using the technology. Instead, they disperse across all services offered. They will not be perceptible by final customers and, even worse, they will go unnoticed by the controlling department, and eventually no incentive for further investment in green technologies will be seen.

The current cost scheme must be therefore changed to direct energy saving to services properly. An initial, rough approach may consist in measuring energy consumption by selected (i.e. “green” devices), and attributing energy cost directly to the appropriate service layer. This requires energy consumption data to be acquired, transmitted, stored and processed. Using cheap, mechanical energy meters means high operational costs; on the other hand, deployment of meters with communication capabilities is more costly.

Another approach is based on traffic measurements. If the relation of power consumption to traffic volume for a device is known, then energy saving can be computed from traffic statistics, like those presented in the figures above. Moreover, in such approach saving can be attributed directly to specific links, or “relations”. The concept of a relation is already present in the current NASK cost model, and the costs of a relation comprise port and link costs, and common costs assigned in a manner similar to what was already presented here.

4.3 Energy efficiency requirements and target energy saving for next generation networks devices

As outlined in section 3.3, the need for greenhouse gas reduction and the energy efficiency requirement of telecom and ISP operators is so urgent and heavy, that the same sustainability of next-generation Internet architecture will be certainly affected. Obviously, energy efficiency and the massive adoption of novel green network devices will play a central role in helping telecom operators (and their customers) to reduce their OPerating EXpenses (OPEX), which today mostly arise from their electric energy bills. At the same time, the massive and fast adoption of such green equipment will certainly depend on the CAPital EXpenses (CAPEX) that will be required for their deployment. Only if the OPEX saving expected from green network technologies and devices will largely overcome CAPEX figures, it is reasonable to suppose a fast deployment of a greener Internet.

In this respect, an explanatory example comes from optic technologies for broadband access (e.g., GPON). Notwithstanding their energy requirements are usually much lower than copper-based protocols, (e.g., DSL), the available bandwidth, and their maturity levels, their deployment is still much limited all around the world. This limited deployment is substantially due to the enormous investments (especially related to excavations), which are required to replace last-mile copper cables. For instance, Telecom Italia owns today more than 100 million km of copper cables. Thus, next research efforts on network energy consumption have to heavily take copper-based broadband access protocols into account.

For similar reasons, clean-slate approaches may be too hazardous in terms of investments, since, by definition, they are not compliant with legacy network technologies, so they may consequently require the replacement of all network devices and links. In this context, the main mission of the ECONET project is to explore approaches and technologies economically viable, and that can be deployed for disruptively reducing the energy consumption of current networks at short-terms. Moreover, clean-slate approaches may certainly have a negative impact on green-house gas emissions, because of the embodied carbon footprint for the massive replacement of network equipment.

For this purpose, the ECONET consortium is considering the internal organization of current networks and devices, trying to identify the most energy starving sources and exploring new solutions for reducing their requirements. Nowadays, it is widely recognised that the sole introduction of low consumption silicon elements may not be sufficient to effectively curb tomorrow’s network energy requirements. Based on this assumption, the ECONET project will investigate, develop and test new capabilities for the Future Internet devices that can enable the efficient management of power consumption in order to strongly reduce the current network energy waste.

The ECONET project will therefore be devoted at re-thinking and re-designing wired network equipment and infrastructures towards more energy-sustainable and eco-friendly technologies and perspectives. The overall idea is to introduce novel green network-specific paradigms and concepts enabling the reduction of energy requirements of wired network equipment by 50% in the short to mid-term (and by 80% in the long run) with respect to the business-as-usual scenario.

To this end, the main challenge will be to design, develop and test novel technologies, integrated control criteria and mechanisms for network equipment enabling energy saving by dynamically adapting network capacities and resources to current traffic loads and user requirements, while ensuring end-to-end Quality of Service. This will be performed at all the various levels of real-world wire-line networks: from the home gateway and access devices, to the ones working at metro, transport and core levels. ECONET solutions for saving energy must obviously target the features and working constraints of each single network layer, but they will need to be led back to two main enabling techniques:

- Power scaling: for tuning the energy consumption with respect to the real utilization of the hardware resource, without losing any network functionality.

- Smart Standby: for intelligently putting the hardware resource in very low energy modes, temporarily losing the largest part of its network functionality, and to intelligently migrate its software/protocol state elsewhere in the network.

Such techniques will be included in the data-plane of all the devices above mentioned, and controlled through control processes, which will estimate the network conditions (e.g., traffic loads, QoS, state of neighbouring devices, etc.) and will decide the new energy-aware configurations in an autonomic or user-guided way.

The use of these two capabilities will provide a composite "energy profile" to next-generation network devices, or their subcomponents (e.g., packet processing engines, link interfaces, etc.). As shown in Figure 85, when devices or parts of them will be not used in the long term (e.g., redundant hardware placed in the network for resilience purposes), they can use smart standby techniques for entering sleeping modes and saving energy. When devices are fully active, they can use power scaling techniques for modulating their capacities with respect to the actual incoming load.

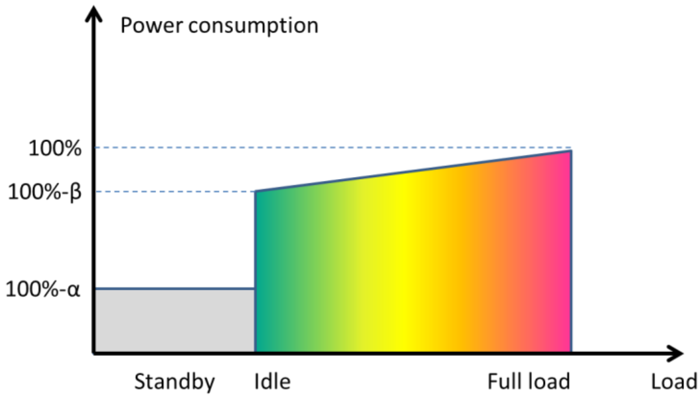

Figure 85: Power scaling and smart standby impact on the "energy profile" of network devices.

For example, as far as the core/transport level is concerned, smart standby primitives can be reasonably used for sleeping redundant hardware devices or components. During low traffic periods, the traffic paths can be modified in order to empty and to sleep further links, hardware modules, or entire network devices. In this respect, the use of state-of-the-art optical technologies (e.g., ROADM, see section 0) can enable next-generation networks to dynamically manage, through reconfigurable optical bypass [115][116][117], the number and the density of silicon-based devices (which usually contain most of the network intelligence, but they are the most energy-starving part of it).

The final objective of the project is to provide of proof-of-concept prototype including advanced power management capabilities, able to halve the energy consumption of an operating network (including the customer premises equipment). An analytical presentation of the energy-aware design space is provided in sections 5.1 and 5.2 for the data and control plane respectively. In summary, in a next-generation network of a Telco or of an ISP, the ECONET technologies will allow:

- *Load dependent power consumption*: Consume energy with respect to the real traffic carried in the network, and not with respect to the deployed link/device capacities.
- *Enabling sleeping modes*: put redundant or hardware components (e.g., links, cards, etc) not used in the long-term (e.g., hours, days) in sleeping modes without causing network instabilities or service interruption.

- *Network support for sleeping states of networked devices:* design new network paradigms for supporting the standby modes of networked devices like PCs, set-top-boxes, etc. without making them "fall off" the network.
- *Apply energy aware traffic engineering techniques*: since networks contain multiple alternative paths to each core network node, specific links could be considered as back-up links and be active only when congestion is present in the network and thus load balancing is considered necessary.
- *Install novel home-access technologies for putting the entire access chain to sleep during user absence:* standby modes of home gateways and DSLAM ports need to be synchronized in order to save energy at the customer home, and in the central office equipment.
- *Apply virtualization techniques in core IP network nodes as well as in data centers*: novel virtualization strategies should be developed in order to provide smart (static or dynamic) consolidation of several physical routers and several physical servers in one virtualization capable router and server, respectively, in order to reduce the number of up and running devices consuming energy.
- *Design autonomic mechanisms for distributed monitoring of energy consumption in network links and nodes:* based on the collected data, energy optimization decisions may be taken for the network.

The rest of this section will discuss the requirements and the research and development actions needed for designing energy-aware optimizations for next-generation network devices, able to disruptively reduce their consumption levels (80% in the long run). The research and development actions here proposed trigger from the energy consumption characterization of today's network devices performed in subsection 4.1, and are expressed per device sub-components in order to better identify where the ECONET research and development effort need to be focused for maximizing its impact.

For each required action, the consortium also identified some performance targets. These targets have been derived by the ECONET consortium with a two-fold approach:

- by starting from today's technological opportunities and first projections on technologies to be exploited in the project (see section 4.4);
- by considering the usual working conditions, device configurations, and traffic load profiles of current and next generation networks (see section 3).

Where possible, the targets have been expressed in a uniform and synthetic way through a small subset of parameters that try to parameterize the device's energy profile above introduced (Figure 85). In more detail, α represents how much energy can be saved if the device (or a sub-component) enters sleeping modes; β represents the energy consumption gain between the state where the device is active but idle and the state when it is loaded at full of its capacity. Additional gain can also be induced in air cooling systems and power supplies thanks to the reduction of operative energy requirements of network specific hardware.

Obviously, these ambitious goals cannot be achieved simply by enabling power management hardware primitives, but they must be supported by network specific mechanisms, protocol extensions, and/or smart procedures to avoid performance/stability drawbacks, and to maximize their impact/exploitation. A first introduction on the specific technological solutions for enabling power management capabilities without causing such drawbacks is provided in section 4.4.

The ECONET consortium strongly believes that if the ECONET technologies will be compliant with these performance targets, a 50% reduction of energy consumption at short-term will be possible, and it will be possible cutting 80% of consumption in the long term with respect to a business as usual scenario through further technological refinements.

4.3.1 Energy saving requirements and targets for network terminations

For energy usage, in addition to the power limits compliance as stated by EU CoC, there are a number of functional requirements the Network Terminations (NT) with routing functionalities should support. An initial list of those requirements is here proposed, most of them derived by contribution given by ECONET partners to the energy task force of the Home Gateway Initiative, further elaborated and customised taking into account of scope, scenarios and objectives of ECONET Consortium. In Table 20, the basic requirements are reported, in terms of functionalities and performance indicators, to be satisfied by a network termination with routing and home gateway functionalities on board.

Table 20: Requirements for network terminations.

<table>
<tr><th>Network Termination component</th><th>Nominal Traffic</th><th>Topology</th><th>Access speed</th><th>Energy usage</th><th>Service constraints for energy usage</th></tr>
<tr><td>LAN</td><td>nominal bitrate of Fast Ethernet (100 Mb/s), Gigabit Ethernet (1000 Mb/s), Wi-Fi 802.11n (300 Mb/s, mean actual value 100 Mb/s)</td><td>NT is the centralized point of connection for all devices in the LAN, and includes both Ethernet switches and Wi-Fi access point</td><td rowspan="3">NT must support nominal bitrate of ADSL, ADSL2+, VDSL2 if the physical WAN interface is copper pair, FTTH GPON if the physical WAN interface is optical</td><td rowspan="3">Power consumption: compliance to EU CoC v4</td><td>Need for real time availability of interfaces for voice services (ETH, Wi-Fi). Small latency needed for other services to pass from low to full power (e.g. IPTV)</td></tr>
<tr><td>FXS and USB</td><td>USB: max. nominal bitrate as per USB specifications (480 Mb/s)
FXS depending on codecs, max. 288 kb/s considering wideband codecs</td><td>USB and FXS are in fact expanding the NT capabilities allowing to connect voice terminals, HDDs, USB keys for various purposes etc. All these devices must be powered by the NT</td><td>Need for real time availability of interfaces for voice services (FXS)</td></tr>
<tr><td>Triple Play</td><td>Combination of traffic related to at least three Eth ports and Wi-Fi connection plus USB stream. Variable depending on service scenarios.</td><td>NT acts as router and centralized connection manager for all devices</td><td>Need for availability to remote access for service purposes (e.g. remote programming of a STB connected to the NT) or remote management (e.g. FW upgrade of NT or home network devices)</td></tr>
</table>

In Figure 86 it is shown the internal architecture of a HG and how it interacts to the external peripherals, based on the Lantiq XWAY VRX288 HG.

Modern system-on-chips provide both digital and analogue integrations enhancing power savings. Lantiq XWAY VRX288 is a universal 6-band DSL System-on-Chip (SoC) in combination with 6-band single chip Analogue Front-end and Line Driver (AFE and LD) (Figure 87) that is specifically designed to be used in HGs. As shown in Figure 87, the SoC supports many communication paths providing full "triple-play" support. On the left side, the AFE interfaces to the DSL.

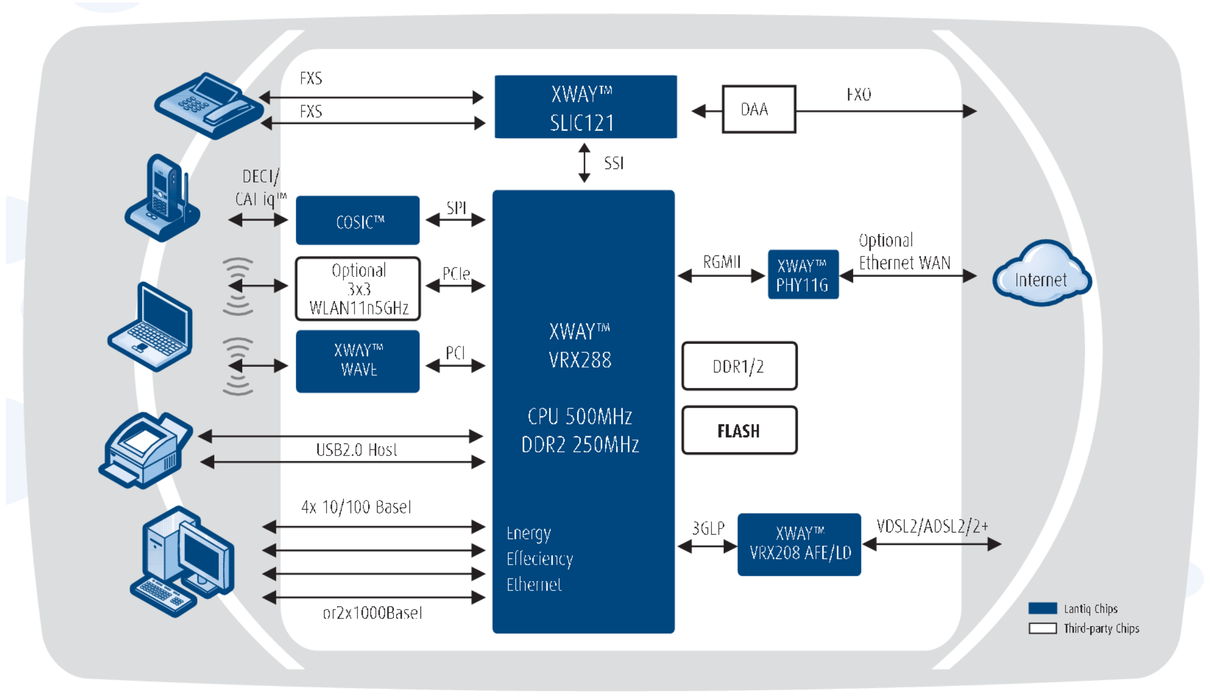

Figure 86: XWAY application example.

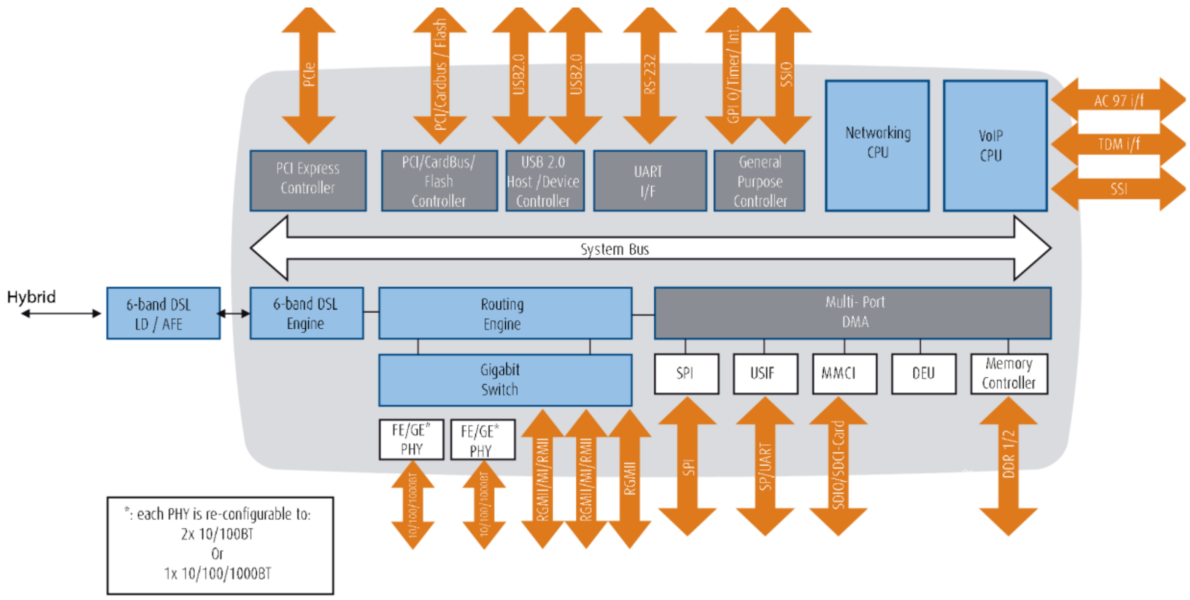

Figure 87: XWAY block diagram.

In general, the NT must support power management functionality, forcing subcomponents to the mode of lowest power consumption that can still support the set of active services. This must be based on automatic activity level detection without requiring any user intervention. Table 22 shows estimations at different action levels for the efficiency that can be reached by the use of the ECONET innovations. For HGs the optimizations are focused at the data plane. The control plane at this network level is supposed to have minimum impact on power losses given its simplicity.

Local interfaces and network termination links have the biggest margin of power saving increase. On the other hand, the packet processing can be less optimized since, as stated in previous sections, the processor is active for every service the HG must support. Air cooling efficiency can be increased only by a more efficient heat sinks design or better low-power small fans. However the passive cooling is preferred in HG devices. The power supply, generally composed of an external part, has a good margin of optimization. The control of the external power supply is suitable to reach 0 W power state. Internal storage capacitors can serve deep low power modes without the use of the external supply.

Currently a DSL interface can be kept idle at about 80% of its peak power consumption and a FXS idle at about 50% of its peak power consumption. The FXS peak is less than 2% of the overall consumption of a HG. DECT cannot be considerably reduced straightaway. WLAN can be reduced to 30% of the peak (Green AP mode) if no clients are available. Ethernet's green mode (EEE) can reduce the power during idle phases to less than 10% of its peak.

In 5 years we do not expect major improvements on EEE and we expect green AP model to reach the same levels of EEE. FXS is expected to show dramatic improvements, but with low overall impact. DECT ULP is also expected to improve a lot its efficiency, so that a single AAA battery can power a DECT ULP link for 10 years. A control plane CPU with its subsystem is currently idle at about 60% of its peak power, because most of the gateway's software needs to be executed even if no traffic is flowing. This is expected to considerably get reduced at less than 30%.

Packet processing engines scale with the number of parallel engines required to process the traffic. In idle mode, one engine is active to serve traffic. By increasing the traffic level, more engines are active - usually 2 to 4. This implies that the idle power consumption of the engines ranges between 25% and 50% of the peak one. DVFS (Dynamic Voltage and Frequency Scaling) techniques can reduce this value by 75%, so in 5 years from now an idle consumption around 6% of the peak one can be expected. GPON interfaces change this figure because they transport more data and demand more engines. G.fast (up to 2 Gbps) DSL would also present the same effect.

Table 21: Device internal level of action for introducing energy-aware optimizations

Resolution action level	Plane
Local interfaces	Data
Network termination link	Data
Packet processing engine	Data

Table 22: Energy-saving targets expressed per resolution level.

Local interfaces	Short Term	Long Term
standby efficiency (α)	40%	50%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	0%	10%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	0%	5%
air cooling (γ)	5%	10%
power supply efficiency (γ)	40%	50%
Network termination link	**Short Term**	**Long Term**
standby efficiency (*α*)	5%	10%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	20%	30%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	10%	15%
air cooling (*γ*)	5%	10%
power supply efficiency (γ)	5%	10%
Packet processing engine	**Short Term**	**Long Term**
standby efficiency (α)	10%	20%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	20%	30%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	15%	20%
air cooling	10%	20%
power supply efficiency (γ)	10%	20%

General power management requirements

The NT must have an on/off switch, except where the NT needs to provide always-on services (e.g. some types of voice services, utility services, security services or guest access). It must be possible for the user to configure a number of periods (by specifying the days of the week when it is applied and the start and end time) during which the NT will operate with only a subset of interfaces active (for example during the night). The mode transition must be completed within 10 seconds of the programmed start or end time. The user has to be able to configure which interfaces are active or inactive during the defined periods.

It must be possible to individually deactivate (completely power down) and activate (power up) LAN side interfaces (e.g. Ethernet port 2 and FXS port 1) via local UI or remote management. In case of deactivation of interfaces, at least one LAN interface must remain active. Powering down all interfaces (so that the user will lock the NT) has not to be supported. It must be also possible to deactivate all power saving mechanisms for troubleshooting via local UI and remote management, as well as to (re)activate the power saving mechanisms that have been deactivated via local UI and remote management.

The NT should provide some visual indication (e.g. LED) of which operational mode the NT is in. The NT should support a LED ECO mode including the provision of a dedicated ECO mode LED. In the case where the NT is in LED ECO mode the NT should provide an alternative way to check the system status (e.g. via local UI). The NT may record the relative power consumption (e.g.

in percentage of the maximum NT power consumption) in the operational mode of all subcomponents over time.

Requirements on subcomponents

Ethernet interfaces

The NT Ethernet interfaces must be compliant with Energy Efficient Ethernet (IEEE 802.3az). The NT must be able to auto-power down a LAN Ethernet port when no electrical signal is detected for a period of 10 seconds so that Ethernet ports with no devices connected will be powered down. The NT must periodically send out link pulses while the port is in auto-power down mode. The NT must detect when an electrical signal is present on a powered down Ethernet port (i.e. if a device gets connected) and activate the port again within 3 seconds. It must be possible to disable the Ethernet auto-power down feature.

CPU

The instantaneous energy usage of the CPU must scale with its actual workload; this implies at least 3 power modes. CPU core(s) must support the low power wait instruction. The CPU core(s) should have power management with, for example, frequency and/or voltage scaling and software policies to dynamically adapt to changes in workload. The CPU subsystem may control the power modes of other subcomponents (e.g. USB, Ethernet, etc.) depending on their load (if not controlled autonomously). In multi-core implementations, unused CPU cores should go to their lowest power mode.

Wi-Fi

The NT must have a physical Wi-Fi on/off button. Operating the Wi-Fi on/off button must override any scheduled actions, i.e.

- if the Wi-Fi interface is off, pushing the Wi-Fi on/off button must power it on until the button is pushed again or the next scheduled Wi-Fi off period begins.
- if the Wi-Fi interface is on, pushing the Wi-Fi on/off button must power it off until it is enabled again by pushing the Wi-Fi on/off button.

It must be possible to configure the Wi-Fi transmit power (e.g. to 4 predefined levels such as 100%, 75%, 50% and 25%) via local UI and remote management. When in low power mode, the Wi-Fi interface must be able to handle user traffic within 3 second by leaving the low power mode. The Wi-Fi interfaces should go into a low power mode (e.g. power down RF chains in MIMO systems) when no user traffic is transmitted for a configurable period (e.g. 2 minutes).

Voice (FXS, DECT)

The voice subcomponent (FXS or DECT) must be put into a low power mode within 10 seconds of an on-hook trigger. The voice subcomponent (FXS or DECT) must be able to provide voice service within 1 second of an off-hook trigger or incoming call. The technique used to detect if a device is connected to the FXS port must not cause connected phones to ring or lights to flash. The NT must periodically check if a phone is connected to a FXS port which is in UNPLUGGED POLLING state and activate that port again on average within 3 seconds after going off hook. The NT may be able to detect a phone or other device connected to a FXS port and disable the port automatically if nothing is connected. The default time interval between checks is 60 minutes. The

DECT interface (if present) must support a base ECO mode to limit the transmit power depending on handset distance.

Power supply

The external power adapter must fulfil all the mandatory requirements defined in ETSI document ES 202 874-1 on Common Power Supply.

4.3.2 Energy saving requirements and targets for access networks

The MSAN (Multi Service Access Node) must support power scaling facilities in order to modulate the consumption in function of the workload. Since MSAN is an "always on" device and the lack of traffic on its ports it is almost improbable, a complete power down of MSAN parts on the data plane is not realistic.

In order to examine energy saving requirements for access networks, we are based on the Lantiq VINAX V3 architecture that achieves low power functions by means of new power saving oriented design. VINAX V3 is a unified platform for ADSL and VDSL applications (Figure 88), that provides the necessary flexibility since it can be re-configured as ADSL2/2+, VDSL2 Profile 8, profile 17a and profile 30a (see Table 23).

Table 23: VDSL profiles characteristics.

Profile	Bandwidth [MHz]	Power [dBm]	Max. Downstream [Mbit/s]
VDSL2 profile 30a	30	+14.5	200
VDSL2 profile 17a	17.664	+14.5	100
VDSL2 profile 8a	8.832	+17.5	50
VDSL2 profile 8b	8.832	+20.5	50
VDSL2 profile 8c	8.5	+11.5	50
VDSL2 profile 8d	8.832	+14.5	50

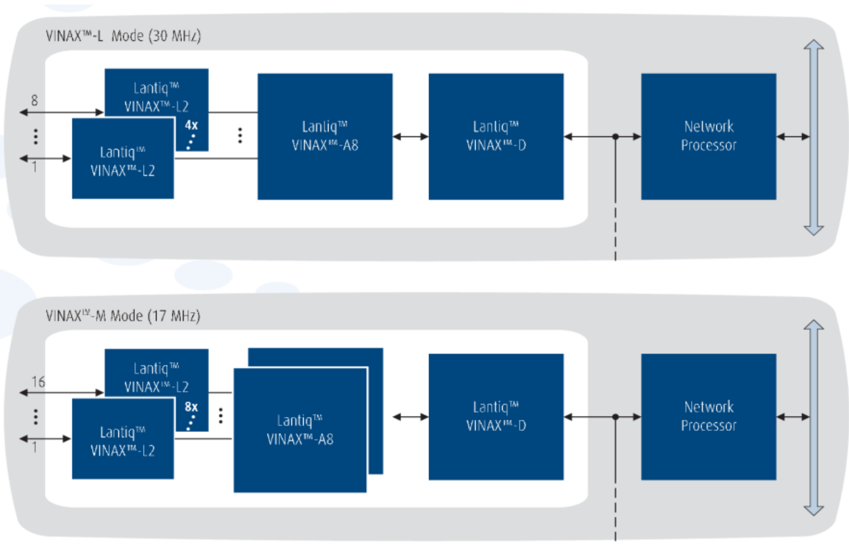

Figure 88: MSAN architecture.

The MSAN energy consumption scales by the number of channels under the different workload. Therefore, the CPU frequency will not be scaled down based on the workload as workload is almost balanced with the scaling number of channels. In order to fulfil all requirements for next generation high definition IPTV as well as VoIP, the VINAX V3 chipset offers full standard compliance to all relevant DSL standards including G.INT (G.999.1), Retransmission (G.998.4), SELT and MELT (G.996.2), and Vectoring (G.993.5).

Table 25 shows energy saving expectations for MSAN line drivers and packet processing engines. In the access scenario, the margin of power saving optimization is increased by the multiple presence of VDSL interfaces through which, a global power saving result in MSAN is achievable.

Table 24: Device internal level of action for introducing energy-aware optimizations.

Resolution action level	Plane
Line drivers	Data
Packet processing engines	Data

Table 25: Energy-saving targets expressed per resolution level.

Line Drivers	Short Term	Long Term
standby efficiency (α)	40%	50%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	20%	30%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	10%	15%
air cooling/power supply efficiency (γ)	5%	10%
Packet processing engines	**Short Term**	**Long Term**
standby efficiency (α)	5%	10%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	20%	30%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	15%	20%
air cooling/power supply efficiency (γ)	15%	30%

Given the low-power architecture of nowadays SoCs, in order to increase the power saving the ITUT G.vector (G.993.5) [123] technique is introduced that allows expanded use of 100 Mb/s DSL. G.vector raises VDSL2 connection speeds up to 100 Mb/s at distances beyond 500 m from the fiber termination point with no transmit power increase and no Shannon-Law violation. G.vector simply removes most of DSL's crosstalk noise, thus providing a very high throughput. A similar principle is used in well-known gigabit Ethernet connections. G.vector is still not regulated by the CoC. It can increase power saving since it permits to maintain the same QoS of the conventional technology at a reduced transmission power [124].

Figure 89 represents a line card equipped with two VDSL digital front-ends and the G.vector engine. The requirements for MSAN include the vectoring function. Although the vectoring engine chip consumes power, the estimations reported in Figure 90 show that the power transmitted over the DSL lines can be largely reduced, because as an effect of the lower crosstalk noise, higher data rates can be achieved at a low power increment.

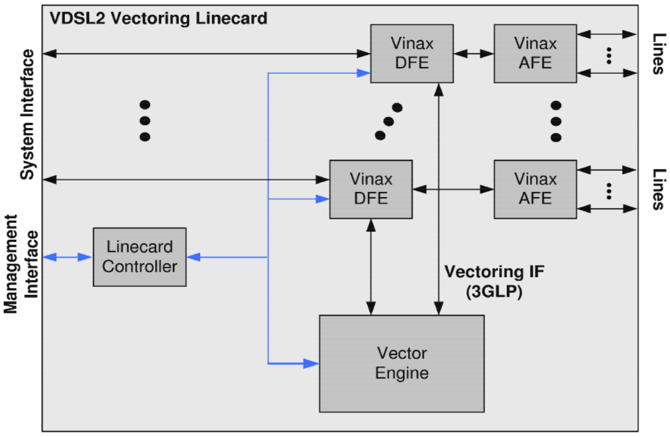

Figure 89: Vectoring Line Card.

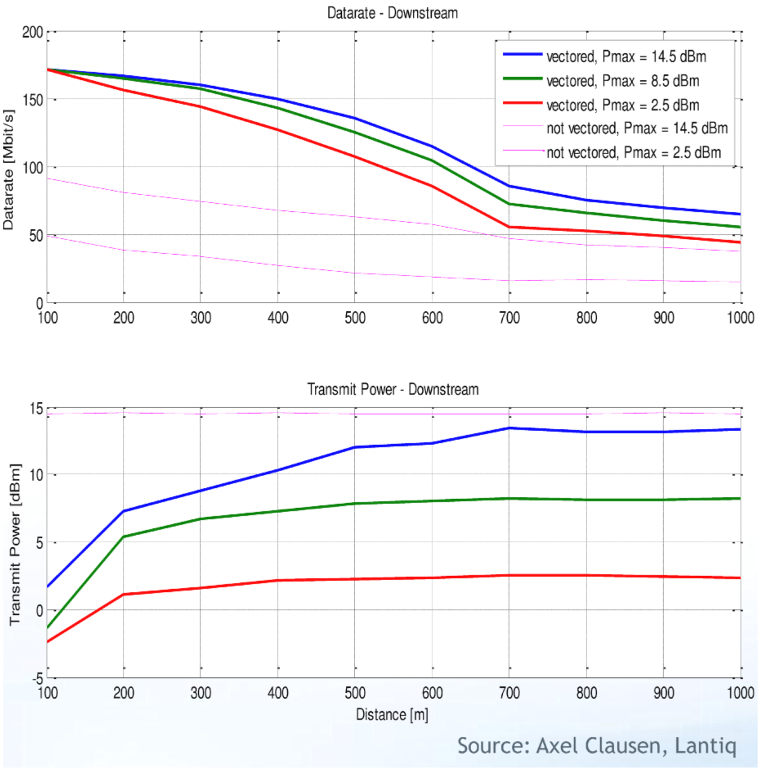

Figure 90: Power saving obtainable by vectoring.

Based on the data in Figure 90, it is possible to see that for a given distance between the MSAN and the NT (for example 500 m), using the vectoring and keeping the same transmit power level (e. g. 2.5 dBm), the data-rate is considerably increased, from 25 Mbit/s to 110 Mbit/s. On the other hand, if the necessity is to save power, it is possible to see that, given a fixed bit rate (e. g. 50 Mbit/s), by the use of vectoring the same performance is obtained spending only 2.5 dBm against the un-vectored 14.5 dBm of power. Besides, the distance is also increased to 900 m against 680 m for the 50 Mbit/s bit-rate example. G.vector, however, has the biggest impact on performance at shortest distance. At larger distance, from 1 km up, it tends to achieve the same trend of the un-vectored case. This is because G.vector is a special equalization on DSL channels, in order to avoid one channel to cover with its power the adjacent ones, keeping them to lower power levels. When the wire attenuation, with the increase of the distance, becomes consistent, the low power levels are more sensitive to it and the same bit-rate at shortest distance cannot be sustained. Accounting both the power consumed by the chip and the power consumed by the line drivers, a case study carried out within Lantiq forecasts a saving of about 2/3 on the total power consumption under normal operations.

4.3.3 Energy saving requirements and targets for metro/transport devices (L2/L2.5 switches)

The power consumption optimization of the L2/L2.5 switches, in next generation metro/transport networks, should be considered a win challenge. Four basic requirements are treated to enhance the power saving of L2/L2.5 switch devices in the ECONET project:

- Lower power components: network elements should use low power components, such as:
  - lower power cooling fan units;
  - low power chips, with dynamic power modulation.
- Power consumption metering of network element equipment devices: each network element equipment device (board, matrix, modules) should integrate circuits to implement the power consumption metering in a static and/or dynamic way.
- On/Off/Standby mode for equipment devices: during periods of low network usage, traffic flow could be optimized in a way that a minimal number of network devices are used. The rest of the devices could be put into a standby/off mode to decrease the energy consumption. Furthermore, dynamic power modulation linked to the actual traffic load could be considered. Optimization should be performed in a coordinated manner, possibly across the network management system or network control plane. In addition to adjusting power to actual traffic (active scaling), a switch system should also support shutting down one or more of its units to achieve substantial power saving in “standby” mode. Standby mode is applied when one or more ports can be shut down due to inactivity of the port. The standby capability can be scaled up in modular switch systems to include shutting down an entire switching module when all of its ports are inactive.
- Advertisement of green Traffic Engineering parameters and energy-aware routing computation: in ECONET the GMPLS signalling should be extended in order to advertise parameters related to the energy consumption for network resources. These “green” parameters should model both fixed (i.e. network equipment power consumption) and dynamic values (i.e. traffic load, number of signal regeneration points on the path). The PCE computation algorithms must be able to take into account constraints related to the attributes describing the energy consumption for network resources. These constrained-based algorithms should be able to optimize the routes according to the combined impact of both resource types.

Next generation switches should also support power saving at several levels. A switch system is usually composed of several HW components and one or more layers of software and firmware. The opportunity to save power depends on the existence of hardware mechanisms that allow the system to change its total power consumption. An additional power saving can come from manipulation of the switch working temperature. Usually a switch system has an active cooling system (fans, water cooling). The cooling system consumes a substantial amount of the switch power. By allowing the switch to operate at a higher temperature, a switch system can invest lower power into its cooling system. An important constraint to increasing the temperature is that the network processor power consumption rises at higher temperature. So, a balance point should be found between the fans’ power consumption and the system power consumption.

However, by actively scaling the power consumption of the HW to fit the actual traffic and lower the heat generated, we can achieve additional power saving from scaling down the cooling system. The following HW mechanisms are the main power consumers on a typical switch system:

- Network Processor – The network processor handles the traffic flowing in and out of the switch system. The processor is composed of data handling units (SERDES), logical and memory units. These units should be designed in a way that allows scaling down their power consumption or even shutting them down when they are not in use.
- Board hardware – Board HW should also be designed to be scaled down when not in use. An example for such HW can be memory arrays that are dynamically used by the network processor.
- Cooling system – The following diagram shows the typical power consumption of a fan unit. As can be see, by allowing the fan to work at a lower speed or a lower air volume a switch system can achieve linear reduction in power consumption.
- CPU – A managed switch usually has an additional management module that has a dedicated CPU to handle chassis and traffic management tasks. As in servers, an additional power reduction can be achieved by scaling down to consumption of the CPU unit according to the actual workload required from the switch.

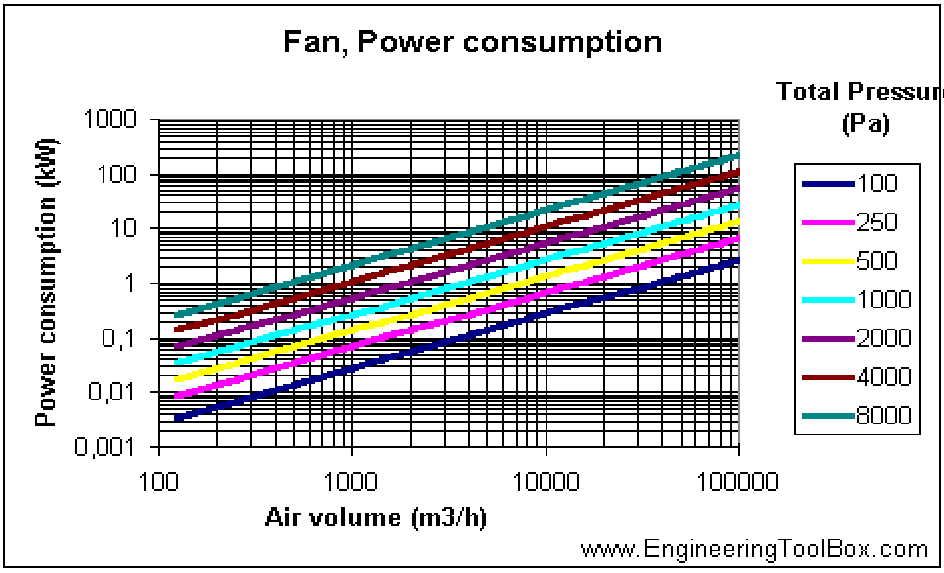

Figure 91: typical power consumption of a fan unit [Source:[11]].

The level of action and energy-saving targets related to the combined energy-aware optimizations introduced is synthesized in the following tables (Table 26 - Table 29). The following resolution is defined for action level: packet processing engines have to involve power scaling and standby on network processor, switch port units and board HW, switch modules have to support standby mechanisms and in the chassis, cooling systems and systems standby mechanisms have to be developed.

[11] http://www.engineeringtoolbox.com/fans-efficiency-power-consumption-d_197.html

In Table 28 and Table 29 the level of action and energy saving targets are specified for the Alcatel Lucent 1850TSS160 switch.

Table 26: Device internal level of action for introducing energy-aware optimizations.

Resolution action level	Plane
Packet processing engine	Data
Module	Control
Chassis	Control

Table 27: Energy-saving targets expressed per resolution level.

Packet processing engine	Short Term	Long Term
standby efficiency (α)	20-30%	40-60%
performance scaling efficiency at load (β)	10-15%	20-30%
Module	**Short Term**	**Long Term**
standby efficiency (α)	50-60%	80-90%
air cooling/power supply efficiency (γ)	10-20%	20-30%
Chassis	**Short Term**	**Long Term**
air cooling/power supply efficiency (γ)	10-20%	20-30%

Table 28: Device internal level of action for introducing energy-aware optimizations.

Resolution action level for Alcatel Lucent 1850TSS160	Plane
Lower power components	Data
Power consumption metering	Data
On/Off/Standby mode	Data
Advertisement of green Traffic Engineering parameters	Control

Table 29: Energy-saving targets expressed per resolution level.

Alcatel Lucent 1850TSS160	Short Term	Long Term
standby efficiency (α)	0%	0%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	0%	60%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	0%	30%
air cooling/power supply efficiency (γ)	30%	50%

4.3.4 Energy saving requirements and targets for core devices (IP router)

Next generation IP routers will have to support different energy saving mechanisms, ranging from power scaling to smart standby. Power scaling primitives will be used for active hardware elements in order to modulate their working speeds with respect to incoming traffic loads, while standby modes will be applied to sub-components not actively working, such as redundant hardware components deployed for network resilience purposes.

In order to maximize the energy-saving impact from the use of these primitives, and given the usual working setup, network scenario, and internal organization of IP core routers (see section 4.1.4), such devices will have to support these primitives at various aggregation levels.

Power scaling and standby mechanisms have to be supported on single link interfaces, on single packet processing engines, on single modules (grouping more interfaces and packet processing engines), on chassis (containing a number of modules), and on single devices' chassis. These capabilities would be appropriately configured by control plane processes and policies depending on the router configuration, service requirements and/or incoming traffic loads. Such processes will span from local optimization policies to network-wide ones (e.g., traffic engineering), and they will be devoted to control the trade-off between energy consumption and network performance.

A summary of actions required for developing next-generation "green" routers is presented in Table 30. Table 31 shows the ECONET targets, expressed as the minimum acceptable level of energy savings with respect to the full load consumption per each level of action resolution (i.e., link, packet processing engines, etc.).

Table 30: Device internal level of action for introducing energy-aware optimizations.

Resolution action level	Plane
Network interface	Data
Packet processing engine	Data
Module	Data
Chassis	Data
Traffic engineering	Control

Table 31: energy-saving targets expressed per resolution level.

Network Interfaces	Short Term	Long Term
standby efficiency (α)	40-70%	90%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	30-40%	50%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	15-30%	25%
air cooling/power supply efficiency (γ)	5%	10%
Packet Processing Engine	**Short Term**	**Long Term**
standby efficiency (α)	80-90%	99%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	45-55%	60%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	20-30%	35%
air cooling/power supply efficiency (γ)	20%	25%

Module	Short Term	Long Term
standby efficiency (α)	50-70%	80%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	40-50%	60%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	20-40%	45%
air cooling/power supply efficiency (γ)	20%	25%
Chassis	**Short Term**	**Long Term**
standby efficiency (α)	50-70%	80%
performance scaling efficiency at 0% load ($\beta_{0\%}$)	40-50%	60%
performance scaling efficiency at 50% load ($\beta_{50\%}$)	20-40%	45%
air cooling/power supply efficiency (γ)	20%	25%
Traffic Engineering (virtualization)	**Short Term**	**Long Term**
standby efficiency (α)	20% additional standby use	30% additional standby use
performance scaling efficiency at 0% load ($\beta_{0\%}$)	-	-
performance scaling efficiency at 50% load ($\beta_{50\%}$)	-	-
air cooling/power supply efficiency (γ)	-	-

4.4 Summary

In section 4, information has been initially provided regarding energy consumption figures of existing devices at the different levels of the network hierarchy, from termination to core IP routers, as well as the requirements set for lowering the energy needs of future devices.

In the home networking part, it has been reported that home gateways are in standby mode for the largest part of the day confirming that access to their features is not continuous. The off period is usually during the night, while energy consumption caused by the power supply is actually the largest part of their overall energy consumption. Additionally, the best-in-class products do not yet represent the maximum efficiency that can be achieved, as the implementation of automatic power management mechanisms is still far from a widespread deployment.

In the access part, it has been shown that the vast majority of the equipment measured has shown a good compliance with the power consumption targets set in the EC Broadband CoC. A huge difference can be depicted between the energy efficiency of broadband equipment related to the same time period but to different technologies. Based on the nominal peak rate per port divided by the related power consumption, ADSL2+ technology is ten times more efficient than the SHDSL one, while GPON technology is more than one hundred times more efficient than the ADSL2+. VDSL2 technology is two to three times more energy efficient than ADSL2+, while the GPON technology is evolving and can already guarantee four times the speed at less than double the power.

In the metro/transport and core networking part, it can be claimed that IP routers have similar architectures with respect to high-end switching systems, since they usually have highly modular and hierarchical architectures. Operational power requirements of IP routers are mainly due to the data-plane that represents the most energy-starving and critical element in the largest part of network devices. Till today, energy efficiency in IP routers substantially means only maximizing the ratio between the system capacity and its power consumption, without taking the possibility of adapting system performance to the real workload into account. For this reason, internal router components do not include power management capabilities like standby, low power idle or performance scaling. The power consumption of a switch is related to the number of ports, the speed / width of data per port, the chassis size, the use of active / passive cables and the existence or not of a management module. Based on measurements on Alcatel-Lucent and Mellanox switches, it has been shown that the major power consumers in switches are the leaf and the spines. Most of the power in these modules is consumed by the network processor, while the next big consumer is the use of active cable and then the cooling system. Furthermore, we have shown that developing new components with very high integration density and low voltage supply leads to a significant reduction in power consumption.

Regarding network cards, the obtained results show that there is a clear indication on the lack of energy-optimization in the design of such boards. Energy consumption in network cards does not depend on the real utilization of the boards, but on the number of attached cables and the speed of links. In higher speeds, the power consumption of a network card is mostly composed of the power used by the network processor on the card itself.

In addition to the energy consumption figures of various devices, specific results have been presented from experiments that were performed to determine how much energy is consumed by typical network devices under various load conditions as well as when they are in the idle state.

Furthermore, three reference network scenarios were presented in detail, focusing on energy consumption on current networks. Energy consumption was measured per network part and the effects of traffic volumes variation (diurnal patterns, traffic volumes per network segment, average loads) on real time energy consumption were examined. The reference scenarios refer to a telecom operator in Italy and two ISP providers in Greece and Poland. In each case, future energy aware technologies that may be used in order to minimize energy consumption have been identified.

Finally, this section includes information regarding the target energy saving for next generation networks devices. Targets have been expressed in a uniform and synthetic way through a small subset of parameters that try to parameterize the device's energy profile. The ECONET consortium strongly believes that if the ECONET technologies will be compliant with these performance targets, a 50% reduction of energy consumption at short-term will be possible, while it will be possible cutting 80% of consumption in the long term with respect to a business as usual scenario through further technological refinements.

5 Technology specifications of energy efficient devices

This section provides a detailed description of the energy aware technologies that will be examined within ECONET on the network device data plane and on the network control plane. Upon the description of these technologies, energy efficient targets and consumption profiles are provided per device element. Monitoring and management issues are also addressed, focusing of the definition of a set of energy aware parameters that will be exposed to control and Operations, Administration, and Maintenance (OAM) applications.

5.1 Future energy-aware technologies on the network device data plane

The ECONET consortium will consider three main approaches for reducing the energy consumption of next-generation networks on the network device data plane, namely re-engineering, power scaling and sleeping/standby. The rest of this sub-section will try to introduce them.

5.1.1 Re-engineering

As previously sketched, re-engineering approaches are devoted to introduce more energy-efficient technologies, and to optimally exploit them inside network equipment architectures. Novel energy-efficient technologies mainly consist of new silicon (e.g., for ASICs, FPGAs, network/packet processors, etc.) and memory technologies (Ternary Content-Addressable Memory (TCAM), etc.) for packet processing engines, and novel media/interface technologies for network links (energy efficient lasers for fibre channels, etc.).

In this respect, the most challenging solution consists in the adoption of pure optical switching architectures [20], which have long been considered the primary candidate for replacing the current electronic based devices. They can potentially provide terabits of bandwidth at much lower power dissipation than current network devices, but their adoption is still far from reality. Current technological problems mainly regard the limited number of ports (less than 100), and the feasibility of suitable buffering schemes.

Focusing on packet processing engines, decreasing feature sizes in semiconductor technology have contributed to performance gains, by allowing higher clock frequencies and design improvements such as increased parallelism. The same technology trends have also allowed for a decrease in voltage that has reduced the power per byte transmitted by half every two years, as suggested by Dennard’s scaling law [28].

Fixed the silicon technology, energy consumption largely depends on the number of gates in the forwarding hardware. In general, every cell, block and gate requires power, making a strong case for structural optimization within the forwarding engine. The number of gates is generally directly proportional to the flexibility and programmability levels of HW engines. Simpler and faster packet forwarding silicon achieves the best energy cost per gigabit, but extreme hardware specialization can lead to limitations in feature sets and updatability of network device functionalities [29].

In this respect, general purpose CPUs typically present the worst case with respect to power efficiency, and the best one to flexibility. Recent multi-core CPUs are designed with 45 to 65 nm CMOS technology and can feature over two billion transistors. They are fully programmable and can perform any packet lookup operation in existence, but this comes at a cost of relatively high power consumption. They can forward several gigabits per second within a power budget of 100-150 W for a high-end CPU.

On the other extreme, fully customized silicon for packet forwarding provides the best energy-efficiency, but very low programmability or flexibility levels. However, the high development cost can be ultimately offset with superior scaling and higher energy efficiency. Custom silicon can currently achieve an energy efficiency level equal to about 100 W. This is almost an order of magnitude better than CPU-based platforms.

Between off-the shelf CPUs and fully custom silicon, there are many intermediate solutions featuring a broad array of price/performance ratios and ranging from packet-optimized network processors to fully configurable CPU arrays, where features and instructions can be added or removed at will.

Starting from these considerations, different researchers [31] [32] recently faced the issue of complexity in network device architectures (and in particular in IP routers) by proposing novel clean state solutions and network architectures for the future Internet. Here, the main idea consists of significantly reducing the functionalities of devices that work at core and transport networks, so that high-end routers and switches may be manufactured with a lower number of HW gates.

In this respect, one of the most promising approaches was suggested by Roberts [31]. He proposed a radical new concept for traffic lookup, which allows next-generation routers forwarding traffic at the flow levels. This approach certainly leads to a more scalable and simple network device architecture with respect to the current ones, which forward traffic at the packet level.

With a similar aim, Baldi and Ofek [32] suggest a synchronous time-based IP switching approach, which allows synchronizing the operation of routers and scheduling traffic in advance. Such approach is based on pipeline forwarding concepts, and specifically time-driven switching and fractional lambda switching. This will allow reducing equipment complexity in terms of header processing, buffer size, switching fabric speedup and memory access bandwidth speedup. They additionally propose to adopt different global time sources, freely available on earth and in space, for driving synchronous operations.

5.1.2 Dynamic adaptation

Dynamic adaptation approaches are aimed at modulating capacities of network device resources (e.g., link bandwidths, computational capacities of packet processing engines, etc.) according to current traffic loads and service requirements. Such approaches are generally founded on two main kinds of power management capabilities provided by the HW level, namely power scaling and idle logic.

Nowadays, the largest part of current network equipment does not include such HW capabilities, but power management is a key feature in today's processors across all market segments, and it is rapidly evolving also in other HW technologies [33] [34]. In detail, power scaling capabilities allow dynamically reducing the working rate of processing engines or of link interfaces. This is usually accomplished by tuning the clock frequency and/or the voltage of processors, or by throttling the CPU clock (i.e., the clock signal is gated or disabled for some number of cycles at regular intervals). For instance, the power consumption of a CMOS-based silicon can be roughly characterized as follows:

$$P = CV^2 f \quad (4)$$

where P is the active power wasted, C the capacitance of CMOS, and V and f are the operating voltage and frequency values, respectively. It is worth noting that V and f are needed to be directly proportional for a correct working of the CMOS silicon. Decreasing the operating frequency and the voltage of a processor, or throttling its clock, obviously allows the reduction of the power consumption and of heat dissipation at the price of slower performance.

On the other hand, idle logic allows reducing power consumption by rapidly turning off sub-components when no activities are performed, and by re-waking them up when the system receives new activities. In detail, wake-up instants may be triggered by external events in a pre-emptive mode (e.g., “wake-on-packet”), and/or by a system internal scheduling process (e.g., the system wakes itself up every certain time periods, and controls if there are new activities to process).

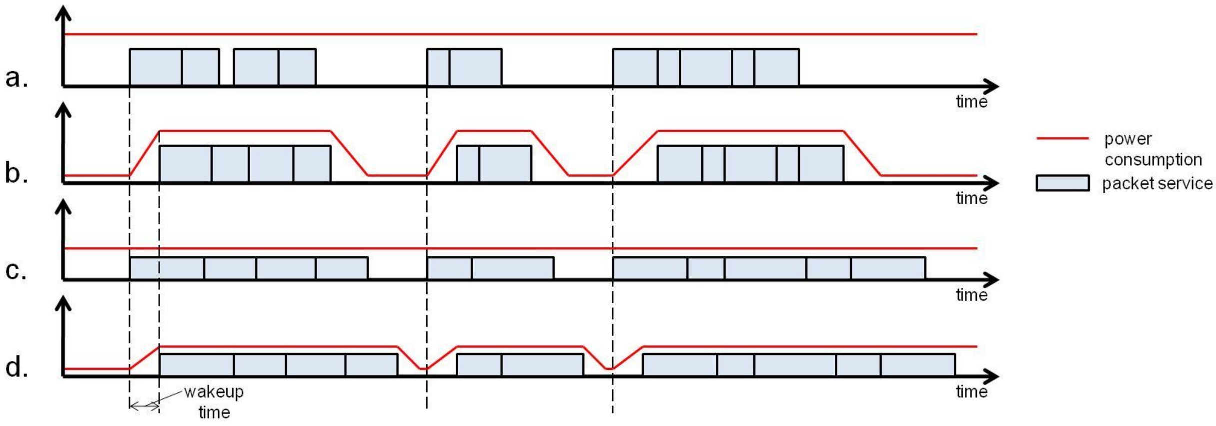

Figure 92: Packet service times and power consumptions in the following cases: (a) no power-aware optimizations, (b) only idle logic, (c) only performance scaling, (d) performance scaling and idle logic.

As in general purpose computing systems, the HW implementation of both idle logic and performance scaling solutions is generally performed by pre-selecting a set of feasible and stable HW configurations, which provide different trade-offs between energy consumption and performance states. For example, different idle states are usually designed by selectively turning off an increasing number of HW sub-elements. On one hand, this leads to reduce the energy consumption during idle times; on the other hand, larger times are needed to wake up all the HW sub-elements.

In a similar way, performance scaling HW support is designed by pre-selecting a set of operating clock frequencies, whose values are sub-multiples of the maximum one and that provide silicon stability.

Both these energy-aware capabilities can be jointly adopted in order to adapt system performance to current workload requirements, and lead to different trade-off between energy consumption and network performance.

As shown in Figure 92, performance scaling (Figure 92-c) obviously causes a stretching of packet service times (i.e., header processing time in a processing engine, or packet transmission time in a link interface), while the sole adoption of idle logic (Figure 92-b) introduces an additional delay in packet service, due to the wake-up times.

Moreover, preliminary studies [35] [36] in this field showed how performance scaling and idle logic work like traffic shaping mechanisms, by causing opposite effects on the traffic burstiness level. The wake-up times in idle logic favour packet grouping, and then an increase in traffic burstiness. On the contrary, service time expansion in performance scaling favours burst untying, and consequently traffic profile smoothing [37].

Finally, as outlined in Figure 92-d, the joint adoption of both energy-aware capabilities may not lead to outstanding energy gains, since performance scaling causes larger packet service times, and consequently shorter idle periods.

However, the energy- and network-aware effectiveness of idle logic and performance scaling (and their possible joint adoption) must be accurately evaluated by taking HW and traffic features and requirements into account. In this respect, it is worth noting that the overall energy saving and the network performance strictly depend on incoming traffic volumes and statistical features (i.e., inter-arrival times, burstiness levels, etc.). For instance, idle logic provides top energy- and network- performance when incoming traffic has a high burstiness level. This is because less active-idle transitions (and wake-up times) are needed, and the HW can remain for longer periods in low consumption state.

Starting from these considerations, Nedevschi [37] firstly proposed to support such energy-aware capabilities (with a special reference to the idle logic) with I/O traffic handling mechanisms, able to shape traffic profiles in order to optimally exploit idle logic and performance scaling. For example, an I/O traffic handling mechanism based on a simple polling policy well suits an optimal use of idle logic.

An optimization policy is generally needed to configure and control the usage of energy-aware capabilities and states with respect to the estimated workload and service requirements. In off-the-shelf computing systems, such optimization policy is usually developed as a SW application, called “governor”.

Regarding the optimization policy, several methods have been proposed in order to estimate the current workload and to optimally control the trade-off between performance and energy consumption in the computing system field. These methods range among predictive techniques [38] and dynamic schemes [39] [40] [41] [42], which were studied for disk drives [43], processors [44] [45] [46], and other components. However, these methods require significant computation to derive the optimal policy and to estimate the current workload, which might not be feasible in all cases.

5.1.3 Sleeping/Standby

Sleeping and standby approaches are founded on power management primitives, which allow devices or part of them turning themselves almost completely off, and entering very low energy states, while all their functionalities are frozen. Thus, sleeping/standby states can be thought of as deeper idle states, characterized by higher energy savings and much larger wake-up times.

The widespread adoption of such kind of energy-aware capabilities is generally hindered by the common aim and design of today’s networking applications and services, which are commonly thought to be fully available all time. In more detail, when a device (or a part of it) goes sleeping, its applications and services stop working and lose their network connectivity. As a result, the sleeping device loses its network “presence,” since it cannot maintain network connectivity, and answer to application/service-specific messages. Moreover, when the device wakes up, it has to re-initialize its applications and services by sending a non-negligible amount of signalling traffic.

In this respect, notwithstanding that PC architectures already include such power management features (allowing desktops and servers to quickly enter sleep and low consumption modes), networking functionalities and applications have often interfered with their effectiveness. This is because of the inability of today's PCs to enter their sleep mode without losing their TCP connections, LAN service broadcasting, etc. This is the main reason why a growing number of networked desktop PCs and servers are continuously left fully powered, even though there is no user demand for their resources most of the time. This is largely because such resources are increasingly shared and must thus be accessible by remote users and other computers 24/7. Moreover, this trend is certainly strengthened by a large part of consumer electronic devices, which are and will be even more “networked” than their PC “relatives”.

Christensen and Nordman directly faced energy efficient enhancements in such kind of scenario [47] [48]. In more detail, their solution to maintain continuous network presence consists of having a network host transfer network presence to a “proxy”, namely Network Connectivity Proxy (NCP), when entering sleep mode [49] [50].

As shown in Figure 93, a NCP is thought to handle ARP, ICMP, DHCP, and other low-level network presence tasks for a network host. A NCP must also be able to maintain TCP connections and UDP data flows and to respond to application messages. Thus, the main objective of such proxy is to respond to “routine” network traffic as the device sleeps, and to wake the device when and only when it is truly necessary.

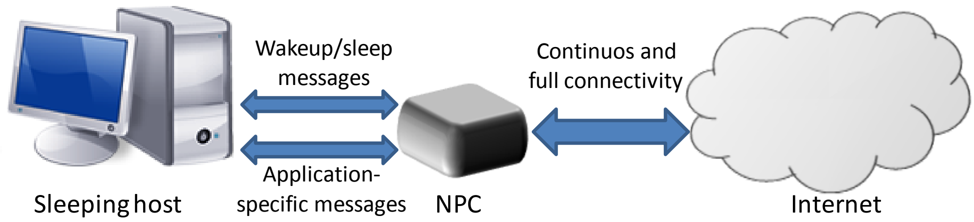

Figure 93: Example of Network Connection Proxy.

In more detail, a NPC and the sleeping device have to exchange two kinds of messages:

- Application-specific: these messages are needed to register sleeping hosts’ applications and services to the NPC. These messages contain the description of application connections, and of application “routine” messages.
- Wakeup/sleep: these messages are needed to trigger the NPC when the host goes sleeping, or to wake-up the host when the NPC receives a message, whose processing directly requires the host.

A NPC can be realized as additional functional block of a network interface, of an Ethernet switch, of a third-party server, or of any kind of device near sleeping hosts.

Energy Efficient Ethernet – IEEE 802.3az

Ethernet is the dominant wireline technology for LANs and is widely used in residences and in commercial buildings. Four different data rates are currently supported in Ethernet using Unshielded Twisted Pair (UTP) as a transmission medium: 10 Mb/s (10BASE-T), 100 Mb/s (100BASE-TX), 1 Gb/s (1000BASE-T), and 10 Gb/s (10GBASE-T). Each data rate was standardized in IEEE 802.3 and uses different modulation and coding requiring different receiver architectures. 1000BASE-T, and especially 10GBASE-T, interfaces are extremely complex mixed signal integrated circuits that include adaptive equalizers, echo and crosstalk cancellers, advanced coding techniques, pre-equalization in the transmitter, etc. The increase in the complexity of the interfaces also implies that more energy will be consumed when the data rate is higher [102].

The approach in EEE is to limit transmission when there is no data to short periodic refresh intervals to maintain alignment between the transmitter and receiver. This is done by introducing the concept of Low Power Idle (LPI), which is used instead of the continuous IDLE signal when there is no data to transmit. LPI defines large periods over which no signal is transmitted and small periods during which a signal is transmitted to refresh the receiver state to align it with current conditions. Large energy savings are obtained when the device spends a significant fraction of the time in the low power mode. Although the savings vary from device to device, the energy consumption when the device is in low power mode can be as low as 10 per cent that of the active mode. During the transitions in and out of low power mode there is significant energy consumption as many elements in the transceiver have to be active. The actual value will depend on the implementation possibly ranging from 50 per cent to 100 per cent of the active mode energy consumption. EEE operation is illustrated in Figure 94 [102].

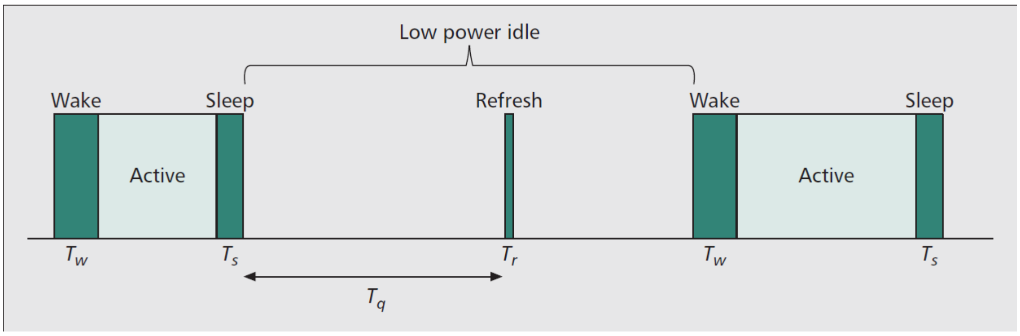

Figure 94: Transitions between the active and low power modes in EEE (Source: [102]).

When packets are being sent, the device is in the active mode, and when no further packets are available for transmission the link may enter the low power (or sleep) mode; the transition to low power mode requires T_s seconds. Once in the low power mode the device only sends signals during short refresh intervals T_r and stays quiet during large intervals T_q. Once packets arrive for transmission, the link is activated again; this wake transition takes T_w seconds. The values of T_w and T_s for 10 Gb/s in IEEE Std 802.3az-2010 are 4.48 µs and 2.88 µs, respectively. In comparison, for 10 Gb/s the packet transmission time, T_{pkt}, is 1.2 µs for a 1500 byte packet (a typical packet size for data transfer using TCP) and 0.0512 µs for a 64-byte packet (a typical packet size for sending TCP ACKs) [102].

Figure 95 displays the energy saving achieved by IEEE 802.3az in home networking over a typical day cycle. Due to the always-on nature of home networking devices utilization, the impact offered by IEEE 802.3az is very large.

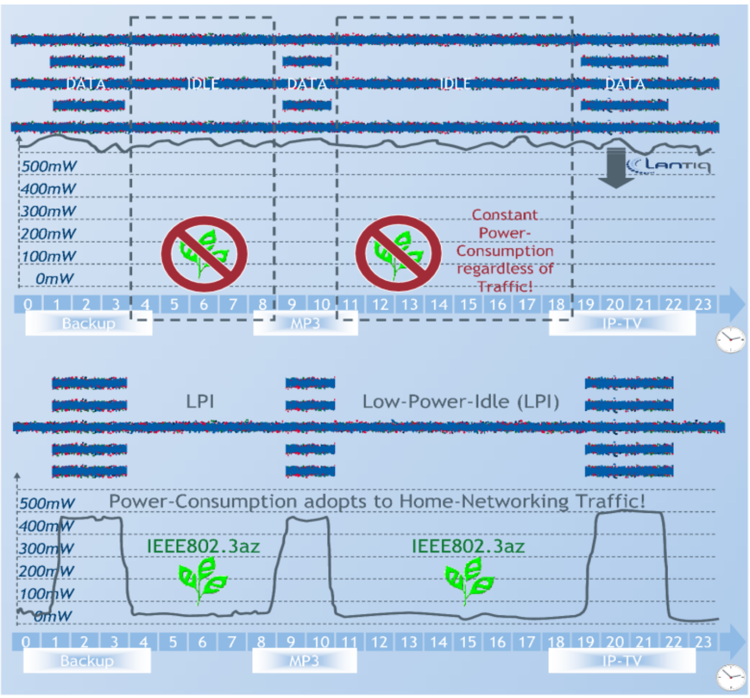

Figure 95: IEEE 802.3az in Home Networking.

Home Gateways, with triple-play features, are required to implement Gigabit links with IEEE 802.3az. IEEE 802.3az defines power states, which the device must switch between in order to save power. 1000BASE-T IEEE 802.3az states are shown in Table 32.

Table 32: 1000BASE-T Low-Power States.

Operating / line state	Description
Active	Legacy operating state where data or idle are transmitted.
Low-power	New operating state used during periods of no data transmission, enabling system power reduction between data bursts.
Sleep	New symbols transmitted to inform the link partner that the local transmitter is entering the low-power state.
Quiet	Transmitter(s) are off.
Refresh	New symbols that are periodically transmitted during the low-power state to allow the link partner to refresh timing and equalization.
Wake	New symbols transmitted to inform the link partner that the local transmitter is returning to the active state. These symbols are transmitted for a sufficient time to allow the link partner to prepare to receive data.

5.2 Future energy-aware mechanisms on the control plane

5.2.1 Energy aware traffic engineering

Network traffic engineering is a method of optimizing the performance of a network mainly considering till now as objectives the user performance and the use of network resources while following a procedure that involves dynamically analysing, predicting and regulating the behaviour of data transmitted over the network. Energy-aware traffic engineering takes traditional network traffic engineering one step ahead by considering and embedding into its objectives the energy consumption of a network in order to achieve the same performance as the energy-oblivious approaches at a lower overall energy cost [68].

Energy-aware traffic engineering is targeted to accomplish networks themselves to become energy-dependent on the traffic served. Networks are designed with two strategic core principles in mind, namely redundancy and bandwidth over-provisioning, which enables them to tolerate traffic variations and faults and work around the Internet's lack of QoS. While these design features do allow network operators to achieve the service level objectives, networks end up being underutilized most of the time, with network equipment typically running at its least energy efficient operating point [71]. The core idea is to dynamically adapt the power consumption of a network by carefully selecting the minimal subset of network elements, which satisfies the current traffic demand. In this approach, one would need to collect information about the current traffic conditions and compute the minimal set of network elements that can accommodate the current traffic, while letting remaining equipment enter the power saving mode. However, putting interfaces on switches or routers to sleep requires additional considerations since it can have serious side effects because of the manner in which various protocols work [72].

In [71] Energy Proportional Networking (EPN) is proposed which is targeted to achieve both optimality and responsiveness by considering a hybrid approach in which: a) as much routing information as possible is pre-computed offline and installed in a small number of routing tables, and b) a simple, scalable online traffic engineering (TE) mechanism is used to deactivate and activate network elements on demand. The high-level overview of this approach to achieving energy-proportionality in networks is shown in Figure 96. Evaluation on ISP and datacentre topologies shows that EPN achieves energy-proportionality while producing substantial savings (up to 42%).

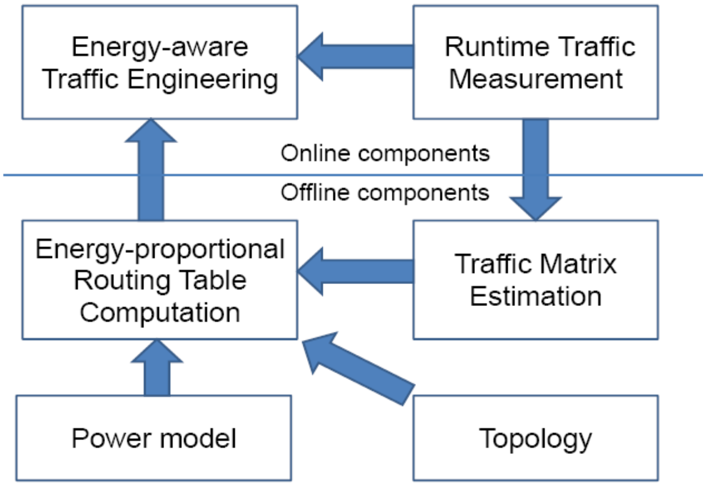

Figure 96: An Energy aware traffic engineering approach (Source: [71]).

GreenTE, another power-aware traffic engineering mechanism, is introduced in [74]. The proposed algorithm aims at maximizing the number of links that can be put to sleep under given performance constraints such as link utilization and packet delay. The algorithm considers that each line-card has a specific base power consumption increasing with the activation of each one of its ports. The proposed algorithm is evaluated considering both commercial and research wide-area network topologies, including the GÉANT network. The results show that the algorithm achieves reduction in line-cards' power consumption by 27% to 42% in conditions where the maximum link utilization is under 50%.

In [73] the authors study energy-aware routing as an optimization problem, evaluating its effects on the network energy consumption and on the network device load as a standard indicator for the QoS performance. Taking as an input the GÉANT network topology several tests have been conducted considering (i) the energy agnostic case where energy consumption of a device is constant, (ii) the realistic case where energy consumption of a device grows linearly between a minimum value above zero corresponding to the idle state and a maximum value corresponding to the full utilization, and (iii) the fully proportional case where energy consumption varies linearly with the device utilization. Results show that energy saving performance strongly depends on both (i) the network topology and traffic conditions and (ii) the device technology, corresponding to different power models. In the realistic case considered, most of the energy saving may be achieved by switching off specific nodes in the network, while switching off links represents a small contribution to the total energy saving. In the fully proportional case, energy saving is achieved through aggregation of traffic to the most energy efficient links. Furthermore, at least for the considered scenarios, it is claimed that achieving energy savings does not necessarily negatively affect network performance although it may raise reliability issues.

In [72] the authors quantify some of the saving that is possible due to inter-packet gaps (i.e., packets not continuously arriving at full speed). Interfaces are put to sleep based on an estimate of the expected inter-arrival time. Saving is highly dependent on the traffic pattern and such opportunistic sleeping intervals might be too short in many cases. In [37] the authors quantify the energy savings when the packets are briefly queued in “upstream” routers, to give “downstream” network elements a chance to sleep longer. According to the presented results the proposed algorithms may even halve energy consumption for lightly utilized networks (10-20%).

In [64], a green provisioning strategy for optical backbone networks is introduced which aims at reducing the operational power following a traffic engineering approach. First, two strategies, namely optical bypass and traffic grooming (i.e., electronically packing multiple sub-wavelength granularity traffic on one lightpath channel) are analysed and evaluated in terms of power efficiency. Then, based on a novel auxiliary graph that captures the operations of signal transmission, an efficient power-aware scheme is devised that searches for an optimal route and a provisioning strategy with minimized operational power. Performance evaluation shows reduced power of the proposed scheme under various network settings, compared with a traffic-grooming approach.

Emerging research approaches to network control, routing and traffic engineering [75] [76] aim at dynamically turning network portions off during light utilization periods, in order to minimize the energy requirements, while meeting the operational constraints and current switching workloads.

Elements in standby (e.g., links or nodes) do literally “fall off” the network, since they are not able to exchange protocol signalling messages to maintain their “network presence” [77]. In core network scenarios, given the features of routing and traffic engineering protocols, the falling off of any elements generally triggers all network nodes to exchange signalling traffic, and to re-converge towards new network logical topologies and/or configurations, causing transitory network instabilities and signalling traffic storms.

Starting from these considerations, the project will investigate a viable approach to introduce standby primitives into next-generation devices, and to smartly support them in order to meet network operational and performance constraints. We will exploit two features already and largely present in today’s networks and devices: the network resource virtualization and the modular architecture of nodes. These features give us the opportunity of using the same base concepts already applied in other fields (e.g., datacentres): decoupling physical elements (e.g., a line-card), which may be put in standby, from their (virtual) functionalities and resources, so that the latter can be migrated towards other active physical elements of the same device.

In a different scenario and with other aims, the idea of virtual router migration was already investigated in [78]. However, in such work the authors suggested the migration of the entire router entity (i.e., its control and data planes) among remote physical platforms. On the contrary, our approach aims at maintaining router entities bound to physical platforms: in this way we can directly control physical nodes, and avoid them to fall off the network.

We will consider a network scenario similar to the state-of-the-art backbone networks deployed by Telcos, where IP nodes have highly modular architectures, and work with a three-layer protocol stack. In more detail, we will consider an IP network (L3) overlaid over a Wavelength Division Multiplexing (WDM) optical network (L1). A Layer 2 (L2) protocol (e.g., the Multi-Protocol Label Switch (MPLS) or the Ethernet protocols) is used to optimally map IP traffic on the physical infrastructure, and to implement value-added network features and services (e.g., Quality of Service (QoS), virtual private networks, mechanisms for fast fault recovery, etc.).

In such environment, physical channels carry multiple “virtual” L2 links (e.g., Label Switching Path (LSP) for MPLS or a Virtual LAN (VLAN) for Ethernet), which directly connect two or more nodes working at L3. Then, each LSP and/or VLAN constitutes a different link at the IP layer.

The path of L2 links on the physical topology is usually determined by using a constrained-based routing algorithm taking into account physical capacity and QoS features. To this purpose, classical IP routing protocols, such as Open Shortest Path First (OSPF), are used within Traffic Engineering (TE) extensions. Moreover, a control protocol, such as Generalized-MPLS (GMPLS), is required to dynamically manage L2 virtual topology.

Regarding devices, we will focus on high-end network routers with modular architectures, composed by a switching matrix and multiple line-cards. Every line-card has one or more Physical interfaces (PHY), and is assumed to include full packet processing capabilities at L2 and L3. As shown in Figure 97.a, each line-card includes multiple PHYs, each one carrying a number of L2 Virtual Links (L2VL). L2VLs are terminated on the line-card itself through virtual network interfaces, called L2 Terminations (L2Ts), which, by definition, are also the network interfaces at layer 3. Thus, IP links are realized by means of L2Ts on two or more nodes.

Given the nature of network protocols, putting entire backbone devices in standby status would not be a practical approach. This is because devices have to maintain their network presence by replying to signalling messages; otherwise they fall out the network and cause a new re-convergence of routing and traffic engineering protocols (e.g., OSPF, GMPLS, etc.). Thus, in order to transparently manage standby primitives and avoid the network falling off, devices must always maintain active control-plane processes, and some connectivity towards other nodes to exchange signalling messages. For instance, if a line-card entered standby status, all packet forwarding operations would stop, and no further signalling messages could be received and/or transmitted by that line-card. Consequently, its PHYs, L2VLs, and L2Ts would fall out the network, as the entire line-card would fault. This triggers fault protection mechanisms for L2VLs to re-converge towards a new L2 topology. Since the terminations of such L2 channel are involved in the topology change, also modifications to the IP logical overlay are highly probable. If the IP logical topology changes, L3 routing protocols must re-converge in their turn, and find new optimal paths.

All this can be summarized in:

- non-negligible amount of signalling traffic across the whole network;
- slow network re-convergence, since both L2 and L3 routing/traffic engineering protocols are involved, and IP protocols generally require long re-convergence times;
- double re-convergence at L2 and L3, which may lead to unwanted traffic paths across the network.

Starting from these considerations, we assume next-generation network devices to have the capability of selectively putting in standby status some of their physical components. In the following we refer to line-cards as the “minimum granularity” building block that can be put in sleeping state, but the approach we propose would be even more beneficial if applied at each line card sub-component (e.g., PHYs, packet processing engines, etc.).

Our idea simply consists of putting to sleep those portions of the device data-plane that are not currently used, like redundant link interfaces, or that are so lightly utilized that their jobs may be temporarily transferred to other active line-cards.

As already introduced, we exploit two features already present in today's networks and devices: the network resource virtualization and the modular architecture of network nodes. These features give us the opportunity of decoupling physical elements, such as line-cards that may be put in standby, from their (virtual) functionalities and resources, so that the latter can be migrated towards other active physical elements of the same device.

In more detail, our idea is manly based on the exploitation of today's L2 protocols for backbone networks (e.g., mainly MPLS and Ethernet), since:

- they are specifically used to manage the virtualization of the physical network infrastructure;
- they already include efficient mechanisms for rapidly moving/migrating L2VLs across the network (e.g., the fault recovery procedures).

In order to avoid unwanted drawbacks in network behaviour, our solution is completely transparent to the L3: IP routing protocols are unaware of network changes and so control message exchange and L3 reconfiguration are avoided.

Our idea consists of making line-cards left active to "cover" sleeping parts, without the device losing any networking resource/functionality. So, before a line-card enters standby status, it has to transfer its resources and activated functionalities to other cards that will remain active.

It is worth noting that such resources and functionalities to be moved are substantially the ones related to all L2VLs and L2Ts carried by the line-card PHYs. As shown in Figure 97.b, we fully exploit the L2 protocols to migrate L2VLs from the line-card entering standby to other line-cards. This obviously requires a new L2VL re-mapping on the physical network topology, since each L2VL has to enter the device from the PHYs of other line-cards. If each L2VL of the sleeping line-card is re-mapped on another active line-card, then the network node sees the same number of L3 interfaces (i.e., the L2Ts), which connect the local router to the same set of IP nodes, as before the L2VL migration. In other words, the full re-mapping at L2 results in a L3 overlay topology substantially identical to the starting one.

Even if no re-convergence of IP routing would be required, standard routers usually consider the L2Ts of re-mapped L2VLs as new network interfaces, since they are allocated on different PHYs and line-cards. A L2T before the migration generally differs from the new one in the interface name/identifier.

Capitalizing on such considerations, our approach simply consists of maintaining the same identification parameters of its old copy in the new L2T.

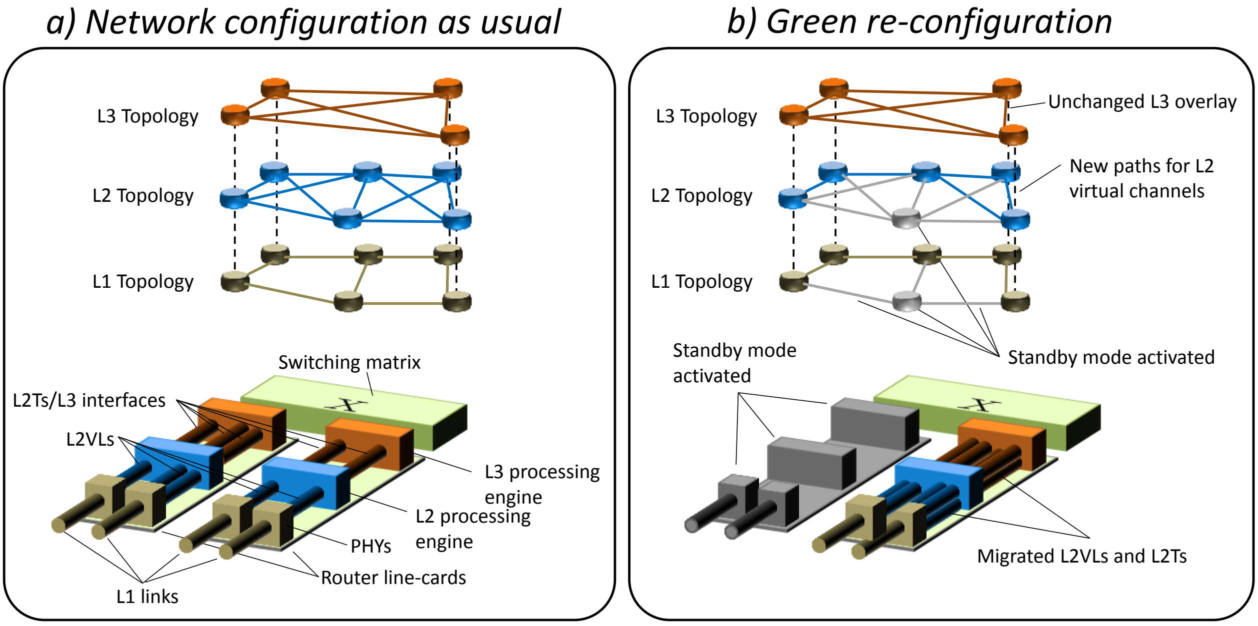

Figure 97: State-of-the-art backbone network and device scenario in subfigure a), and, in subfigure b), the proposed approach to enable network devices to selectively sleep their components. The approach is fully based on network re-configuration at L2, and aims at managing standby primitives in a transparent way with respect to the L3 overlay.

5.2.2 Resource allocation & Algorithmic efficiency

Efforts regarding reducing energy costs are concentrated either at (a) reducing energy consumption (number of kWh) or/and (b) at reducing the cost of consumed energy (cost per kWh). In both cases efficient resource allocation strategies may play an important role. For example under the (a) case, followed consolidation strategies within the cloud data centre [96] or virtual network embedding algorithms in virtualized infrastructures [89] have great impact on the overall energy consumption. Under case (b), the selection of building data centres where average energy cost is relatively low, or even developing strategies where traffic (for a service) is routed to those data centres (from a geographically dispersed infrastructure of datacentres) with the minimum cost for energy in the period under consideration [97], proves critical for the overall energy cost. The latter energy-aware routing strategies could also be used to direct traffic to rely on energy that is produced in a more environmentally friendly way, or even to cut energy usage, by routing traffic away from data centres experiencing warm weather, thereby allowing equipment to be shut down to avoid energy consumption due to cooling requirements.

Furthermore, a recent study [90], [98] has revealed that a single Google search releases 5-10 grams of carbon dioxide (CO_2). Google’s answer refers to 0.2 grams of CO_2 for a typical search. Google's infrastructure replicates queries across multiple servers, which then compete to provide the fastest answer to one’s query. Considering a simple website browsing, total energy consumption comes from the visitor’s machine, the network infrastructure used for the transmission and the data centre providing the website. It is estimated that just browsing a basic website contributes to the generation of about 20mg of CO_2 per second while complex websites with rich animations and video can be responsible for the emission of CO_2 at up to 300 mg per second. Irrespectively of the accuracy of the reported numbers, the fact is that such studies have motivated researchers towards considering the connection between energy consumption-environmental impact and algorithmic efficiency; building more efficient algorithms requiring less computing resources benefits energy efficiency and thus has lower environmental impact.

5.2.3 Virtualization techniques

Virtualization is the technology that combines or divides resources to present one or many operating environments [68], [80], [81]. Virtualization examples are the creation of virtual versions of operating systems, servers, storage devices or network resources.

Server virtualization

Platform (or Server) virtualization refers to the division of a hardware platform (host) into several isolated virtual environments (aka virtual machines, guests, instances) through the use of host software. The main benefits coming from utilizing platform virtualization for implementing a server consolidation strategy may summarize into the following:

a) Reduced hardware maintenance costs from retaining a lower number of physical servers.
b) Efficient utilization of hardware resources. Reduced total cost of ownership (TOI) and faster return of investment (ROI).
c) Reduced administrative costs when having different applications/services running isolated in their own virtual servers by preventing impact of one application to the other when upgrades or changes are made.
d) High diversity of OS hosted at a single platform.
e) Reduced server deployment by having a standard virtual server build that can be easily duplicated.
f) Increased flexibility / agility exploiting the fact that virtual machines may easily migrate to another host server either to illustrate reliability, quality of service related or other business strategies (such as strategies to reduce required operating energy cost).
g) Increased space utilization efficiency in the data centre.
h) Highly reduced energy consumption coming from operating less physical servers while requiring less energy for cooling at the data centre.

We further justify the reduced energy consumption argument which is also highly associated to efficient utilization argument. According to [82] even an energy-efficient server consumes almost 50% of full power when idle (see Figure 98) while it has been observed that servers are rarely completely idle and seldom operate near their maximum utilization, instead operating most of the time at between 10 and 50 per cent of their maximum utilization levels (see Figure 99). Thus, it is clear that server consolidation may provide in the majority of cases significant reduction in energy consumption. Furthermore, according to [83], [84] while in theory the best cooling systems require about 0.3 watts to cool 1 watt of equipment, real-world results from actual data centres, including all parts of the cooling cycle, run between 0.5 and 1 watt of cooling power for each watt of dissipation. According to a simplified example presented in [92] a single high-performance 300W server consumes 2628 kWh of the energy per year, with an additional 748 kWh in cooling, considering that a 12000 BTU air-conditioning with common energy efficiency rating of 12 is used for cooling (cools 12000 British Thermal Units (BTU) for 1 kWh; 1 kWh =3414 BTU). Besides, this also causes a significant adverse impact on the environment in terms of CO_2 emissions that should be taken under serious consideration. So, taking into account that power and cooling are locked together makes the benefits of server consolidation even higher regarding the total required energy for operation.

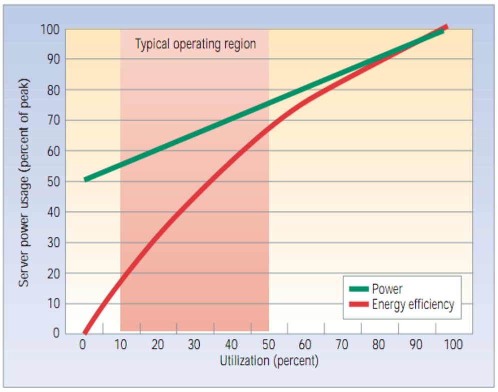

Figure 98: Server power usage and energy efficiency at varying utilization levels, from idle to peak performance (Source: [82]).

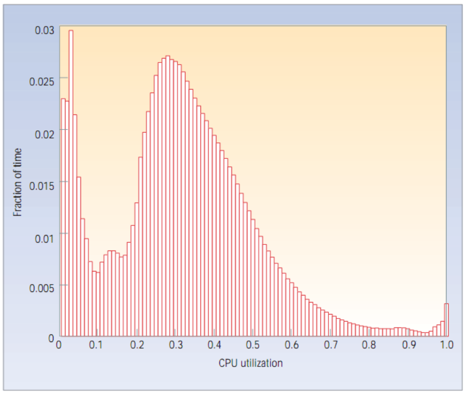

Figure 99: Average CPU utilization of more than 5000 servers during a six-month period (Source: [82]).

Live migration

A feature considered as mature by now, for almost any hypervisor product, is the live migration feature [85]. Live migration allows a server administrator to move a running virtual machine or application between different physical machines without disconnecting the client or application. This feature has facilitated a number of proposals targeted to improve the energy efficiency when operating a series of servers illustrating a number of services. A recent proposal, LiteGreen [86] saves desktop energy by migrating the idle desktops in enterprise environments (running in VMs) to a central server. Thus, the user's desktop environment is "always on", maintaining its network presence fully even when the user's physical desktop machine is switched off and thereby saving energy. This seamless operation allows LiteGreen to save energy during short idle periods as well (e.g., coffee breaks), which is shown to be significant according to the authors' analysis of over 65,000 hours of data gathered from 120 desktop machines. Their findings from a small-scale deployment comprising over 3200 user-hours of the system as well as from laboratory experiments and simulation analysis indicate energy savings of 72-74% with LiteGreen compared to 32% with existing Windows and manual power management.

Another close in philosophy approach that does not however utilize live migration of a whole virtual machine is SleepServer [87], according to which lightweight virtual images of sleeping PCs are created, and these pared down images maintain connectivity and respond to applications, such as Voice over IP services, on behalf of the sleeping PCs. Each virtual PC image can also enable remote access to the sleeping PC it represents via protocols such as Remote Desktop, VNC and encrypted connections using SSH. SleepServer seamlessly wakes up the physical PC when its owner tries to connect remotely into the machine from home, thus enabling a remote connection without requiring the PC to remain on for the entire night or weekend. SleepServer will also wake up the physical PC when the user needs to remotely access stored files and media. According to the measurements presented it is shown that significant energy savings maybe achieved for PCs ranging from 60%-80%, depending on their use model.

A critical factor for both LiteGreen and SleepServer and same philosophy approaches is the time required to wake up and make the original working PC functional (time to resume). Both approaches in the average case require about 20 seconds resuming the original working environment leaving space for optimization. Both works take one step further prior work on putting physical machines into sleep while retaining their network connection (like SleepProxy [91]) by achieving also to perform tasks on behalf of the sleeping machines.

Regarding related standardization efforts Distributed Management Task Force (DTMF) Virtualization Management (VMAN) standard includes a set of specifications that address the management lifecycle of a virtual environment. VMAN's Open Virtualization Format (OVF) specification [88] specifies an open platform-independent, efficient, extensible, and open packaging and distribution format for virtual machines which facilitates the mobility of virtual machines.

Furthermore, the live migration feature has motivated a series of other energy optimizations in virtualized Cloud data centres such as achieving energy saving by continuous consolidation of VMs according to current utilization of resources, virtual network topologies established between VMs and thermal state of computing nodes [89].

Network virtualization

Network virtualization refers to the virtualization of network elements building a network. Network consolidation strategies prove to provide significant benefits. The core building block of network consolidation is the virtual router. A virtual router is defined to be an isolated logical router process and just like a physical router, forwards packets between dedicated interfaces using routing protocol processing while having dedicated interfaces, routing tables, routing processes, radius servers, and even administrative accesses.

There are two main techniques for creating virtualized router entities as defined by their physical and operational characteristics. A Hardware-Isolated Virtual Router (HVR) has hardware-based resource isolation between routing entities, whereas a Software-Isolated Virtual Router (SVR) comprises software-based resource isolation between routing entities [93]. In Table 33 a comparison of the characteristics of the two architectures is presented. As it is also presented in Figure 100 in an SVR there is a constant contention for resources since control plane and data plane resources are shared among the illustrated virtual entities, requiring careful planning and policy construction to ensure that one SVR does not negatively impact others in the chassis. On the contrary in an HVR there are dedicated control plane and data plane resources allowing HVRs to be upgraded or downgraded individually without affecting the operation of the others. Furthermore, HVR solution is highly scalable allowing also illustrating SVRs within each HVR. For example with Juniper's JCS1200, one can build up to 12 hardware-based virtualized routers (six if redundant), each of which supports up to 16 software-based logical routers, for a total of 192 logical routers (96 if redundant) in one chassis [94].

In Figure 101, an example of vertical and horizontal consolidation at a PoP is presented. Horizontal consolidation is the combination of multiple platform functions such as provider (P) and provider edge (PE), or core and edge, into a single platform. Vertical consolidation is the combination of multiple single-purpose devices, such as multiple PE routers, into a single platform. Together, horizontal and vertical consolidation of network elements can reduce the complexity and cost of PoP architectures [94]. Either we are considering consolidation of network elements at the PoP or at the data centre, the benefits are similar to those of platform virtualization regarding energy consumption [90], [70].

Table 33: Comparison of virtualized routing architectures (Source: [93]).

Category	Hardware-Isolated Virtual Router	Software-Isolated Virtual Router
Control plane resources (CPU, memory)	Dedicated	Shared
Data plane resources (forwarding engine, queues)	Dedicated	Shared
Chassis resources (power supplies, blowers, fabric)	Shared	Shared
Management, configuration	Dedicated	Typically shared, but varies depending on degree of virtualization
Connections between virtualized routing entities	Typically external	Typically internal, but possibly external
Per-chassis scalability (routing adjacencies, prefixes)	Increased with additional logical routers	Unaffected by additional virtual routers

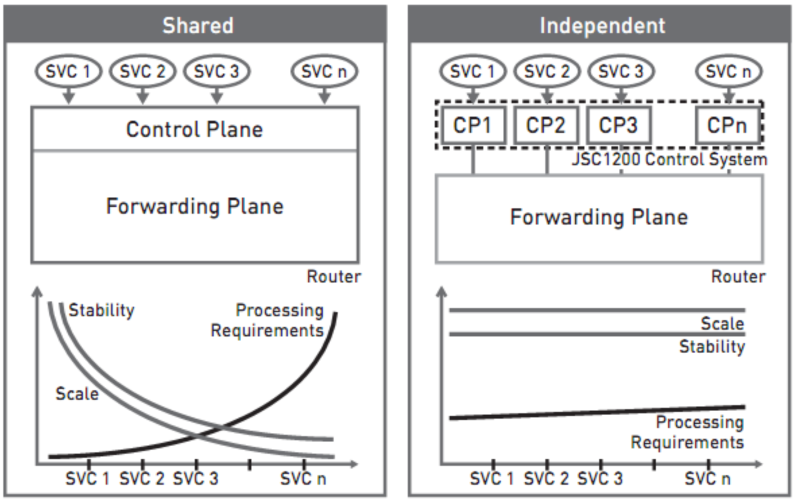

Figure 100: Shared and independent control panel router virtualization architectures (Source: [85]).

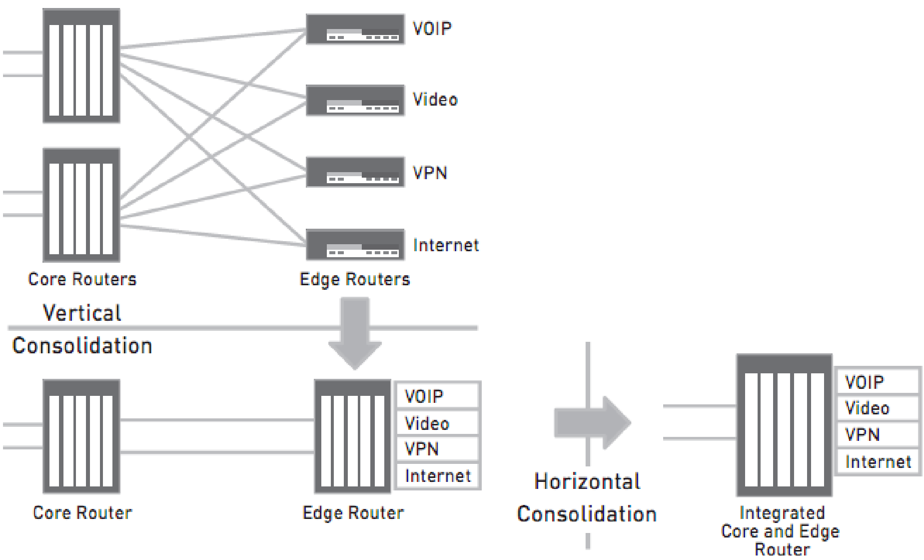

Figure 101: Horizontal and Vertical consolidation (Source: [94]).

Virtual infrastructures

A worth mentioning exploitation of the virtualization technology to offer virtualized infrastructures [96] to researchers experimenting on future internet technologies, is done by the European project FEDERICA [95] which has built an infrastructure based on the multi-domain European National Research and Education Networks (NRENs) and the GÉANT2 backbone. FEDERICA resources are hosted at Points of Presence (PoPs) of participating NRENs (Figure 102). These resources are circuits, switches and nodes capable of virtualization. FEDERICA uses virtualisation in computing and network systems, to create a technology agnostic and neutral infrastructure. It creates “slices” from this substrate, which are a set of virtual network and computing resources according to the user’s request. Virtual slices of FEDERICA’s infrastructure may be created, allocated and used simultaneously by researchers for testing, even with disruptive experiments, within a large production substrate. The researchers have full control of the allocated virtual nodes and network in their slice and can access specific network monitoring information. The infrastructure capabilities allow a much faster research and experiment cycle.

Although such an approach was not initially targeted to energy efficiency when implementing network infrastructures, the energy saving is quite obvious, since providing an Infrastructure as a Service (IaaS) provides aggregate energy saving from both network and platform consolidation compared to physically implementing each distinct hosted infrastructure.

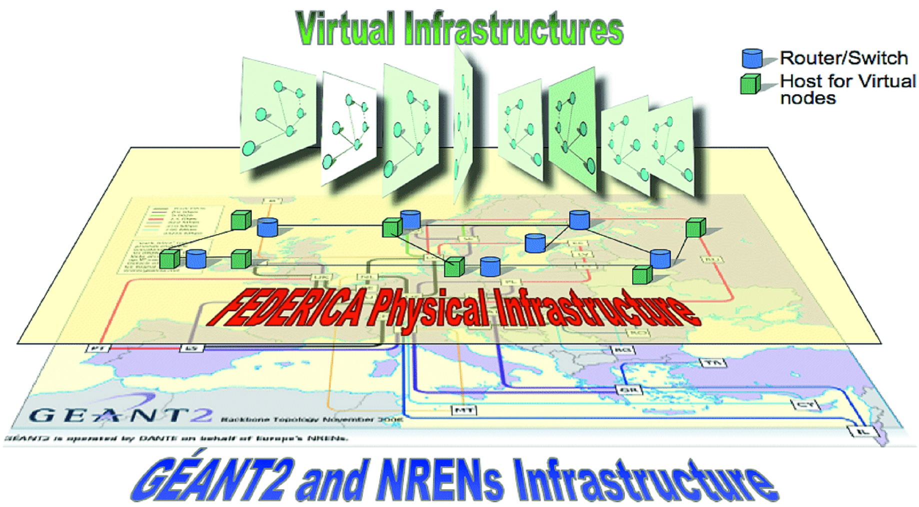

Figure 102: The FEDERICA e-infrastructure.

5.3 Energy consumption profiles for each network device element

As stated in previous sections, there are many technologies that can be applied to different levels of the network in order to achieve meaningful reduction in power consumption. However not all technologies are applicable to all the devices in the network. This section specifies, for the network devices in the ECONET project, which will be the technologies that can be implemented and what will be the power benefit expected from applying these technologies.

5.3.1 Energy consumption profiles of switch systems

Optical networks are high-capacity telecommunications networks based on optical and electronic technologies and components that provide routing, grooming, and restoration at the wavelength level as well as wavelength-based services. Optical transport networks are evolving quickly to adapt to the ever-increasing demands of telecommunication needs, most noticeably witnessed by the explosive growth in transmission capacity demand. The channel data rate in modern optical transport networks is migrating from 2.5 to 10 Gb/s. 40 Gb/s per-channel data rate is on the horizon for commercial identification and deployment, and 100 Gb/s per-channel data rate has been accepted by standards bodies for the next-generation Ethernet. Further evolution to 400 and 1000 Gb/s is also expected within the next three years.

A switch system is able to support multiple ports transmitting data at various speed and width. By applying smart adaptation of power consumption of each port to the actual traffic the switch serves, we can accomplish power reduction up to 50% of the maximum power consumed by today's switching systems. In Figure 103, we can asset the power benefit of applying width/speed reduction techniques. We can see that for a switch operating with 36 ports in active state we can achieve up to 60% power reduction by adjusting the port state to the actual traffic that flows through it.

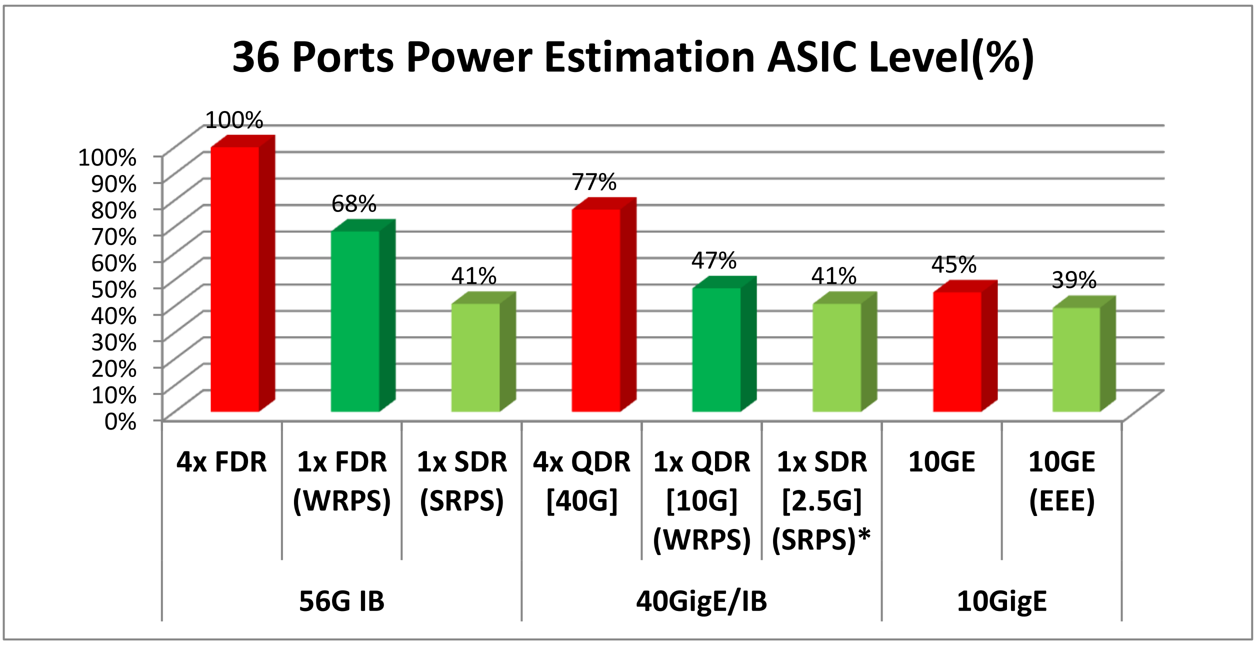

Figure 103: Power benefit of applying width/speed reduction techniques.

In addition, we can see that shutting down a port will cause an immediate reduction in power consumption. By applying smart port shut down techniques, we can achieve power reduction of up to 50%.

A switch system can be composed of one unit that contains a fixed amount of ports or can be a modular system with several independent switching modules. In the case of a modular switch another technique can be implemented, by shifting traffic from one switch unit to the other in order to be able to shut down one or more of the switch modules and save power by doing so.

We can examine the MLX modular switch with 648 ports that is composed of 36 front switching modules called leafs and 18 back connectivity modules called spines. As seen in Figure 104, closure of unused ports or even shutting down a whole unit, when traffic is significantly low, will have an immediate and significant impact per each switch module as on the total power consumption of the entire modular switch. As we have seen in previous sections, this is not the case in today's switches.

Another significant power consumer in a switch system is the cooling system. As seen in previous section the power consumption can take up to 20% of the total power of the switch. The power consumption of a fan is mostly linear, with saturation points at both ends of the working spectrum.

By applying smart cooling techniques we can reduce the average working point of the fans to 50% of max speed and by doing so achieve additional 10% reduction in total power consumption. However, operating at lower fan speed usually means that we will need to raise the working temperature of the switch module and by doing so increase the power consumption of the network processor. It is our job in this project to find a working point in which the power benefit from operating at lower fan speed is larger than the additional consumption due to rising of ambient temperature inside the switch.

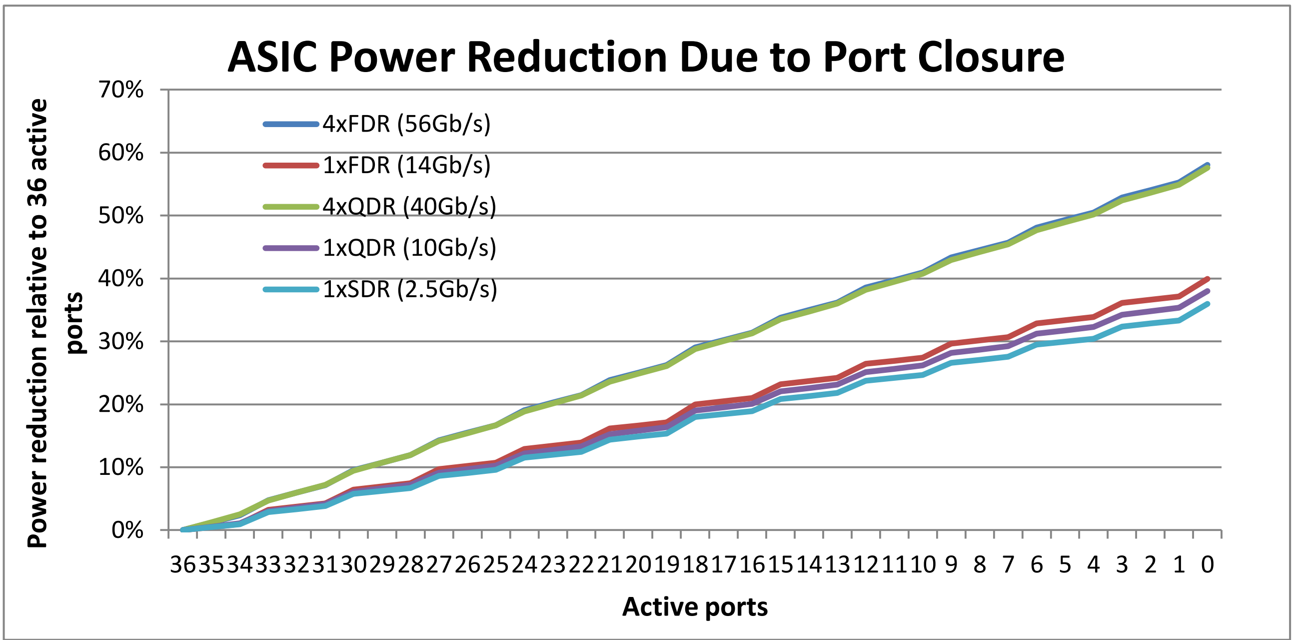

Figure 104: ASIC power reduction due to port closure on 36 ports switch module.

Regarding the 1850 Transport Service Switch, the power figures in a full configuration, are described in the Table 34:

Table 34: Power figures of Alcatel 1850 Transport Service Switch.

	Conf PKT only (70 Gbe+ 8x10Gbe)		
	Max packet configuration with 2KW DCDC Step up (TRU)		
# of boards		Worst	Typical
1	Shelf & Matrix incl. FAN and control platform (full protection)	266	216
8	1X10GE PACKET MODULE SYNCH ETH (1S)	100	80
7	10X1GE PACKET MODULE SYNCH ETH (1S)	100	90
	MULTISERVICE PACKET MODULE (1S)	110	110
	TOTAL Power Consumption	1766	1486
	Expected improvments due to technology (ASIC/FPGA) 20-30%		
	Basic measurement and control tools will open the platform to futher enhancements		

The reference system is a modular packet switch, composed of eight 1x10 GigaBit Ethernet modules and seven 10x1 GigaBit Ethernet modules. By applying technology improvements such as low power components, low power chips with dynamic power modulation and lower power cooling fan units, the maximum expected energy improvement is about 30%. An additional significant power improvement in the proposed packet switch is the integration of power management and control tools that will have additional impacts on the total power consumption of the whole system. Continuous power monitoring and close chain control of the network resources are the reference models of the management and control tools, planned for further power consumption reduction.

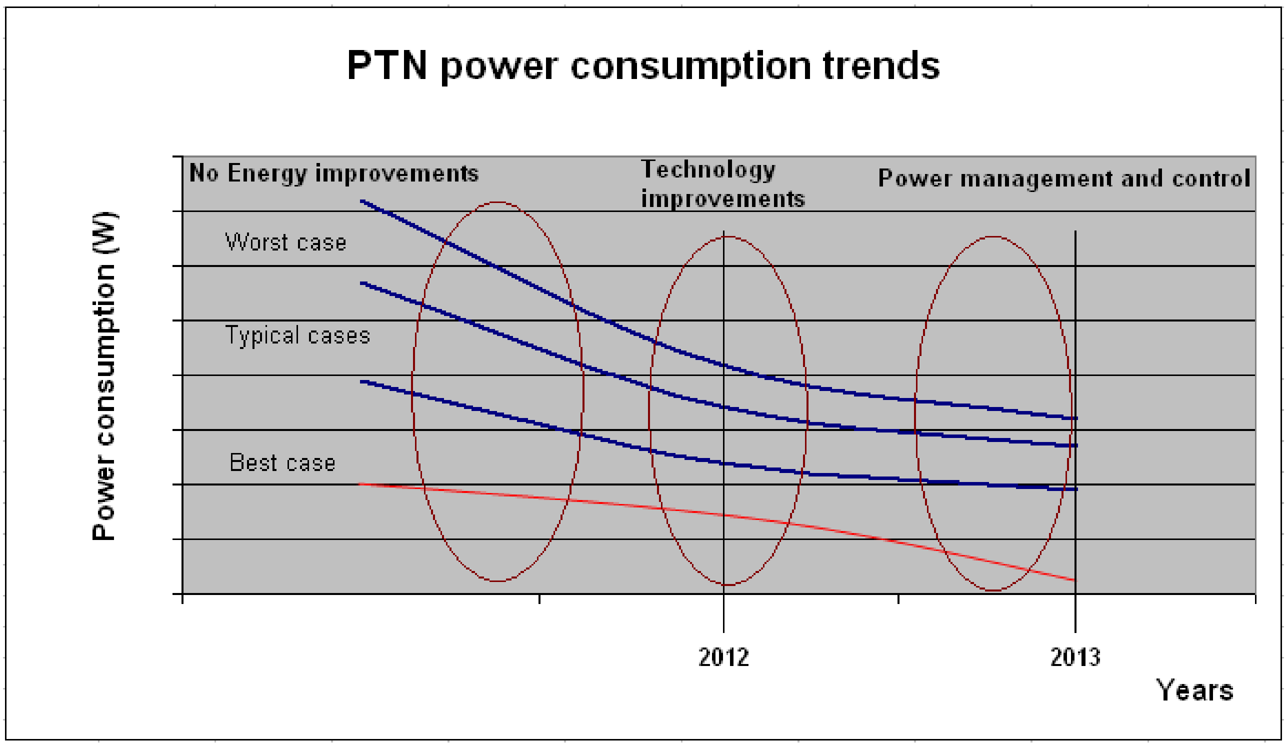

Figure 105: Power consumption reduction of the packet transport Alcatel switch.

Figure 105 shows the power consumption reduction of the packet transport Alcatel switch, planned in the next years: the first area shows the power consumption of the Alcatel switch without energy improvements (in three different cases: best (red line), worst and typical (blue lines)), the second area shows the power consumption of the same switch including energy technology improvements (low power components, chips, etc.) and finally, the third area shows the power consumption of the same switch with power management and control additional improvements.

5.3.2 Energy consumption profiles of VDSL home gateways

For Home Gateways (HGs), it is possible to achieve power saving turning to standby modes unused interface functions during the fruition of one particular HG service, for example a VoIP call. Let us consider the scenario that is shown in Figure 106.

FXS, CPU, memory and xDSL front-end are used in a VoIP call, so it is possible to turn to low power modes the unused interfaces, like DECT, WLAN, Ethernet and USB, naturally in the case they have not to serve other users. This approach takes the advantage of unused or partially used blocks. The result is that is possible to lower more the power consumption and have more margins over the European CoC power targets, as it is shown in Figure 107 (the optimized HG is pointed by the arrow).

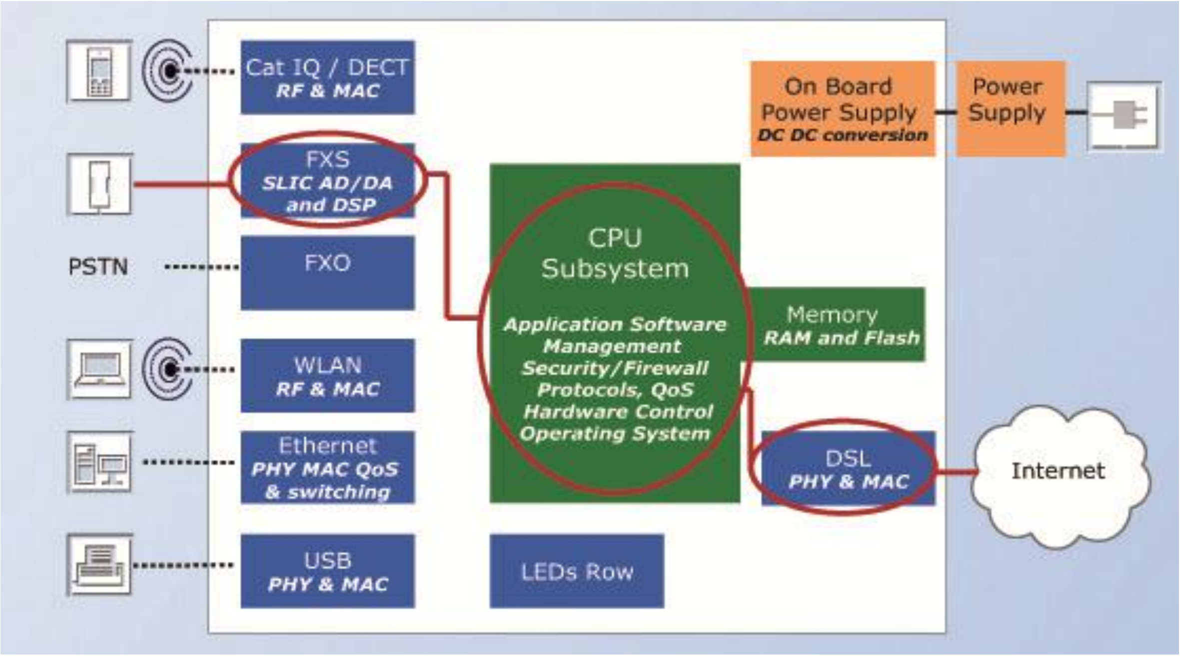

Figure 106: HG functional units in use during a VoIP call.

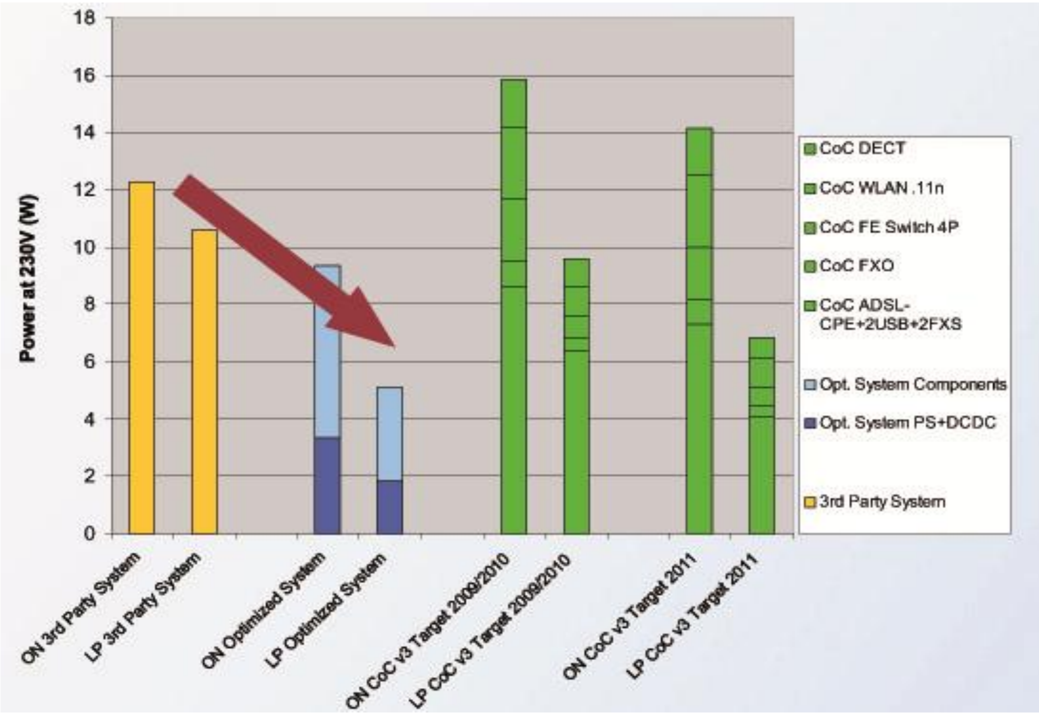

Figure 107: Power consumption comparison of 3rd party system and optimized system against European CoC.

However, Figure 108 suggests that the power supply system is the more power hungry part of the HG. Naturally it has to be also optimized. In particular the external power supply has an efficiency of 55% and the internal DC-DC converter of 70%. By the use of optimized power designs it is possible to achieve efficiency of 80% for the external power supply and of 80% for the internal DC-DC converter, thus having a total 40% reduction in power losses compared to old designs (see Figure 107).

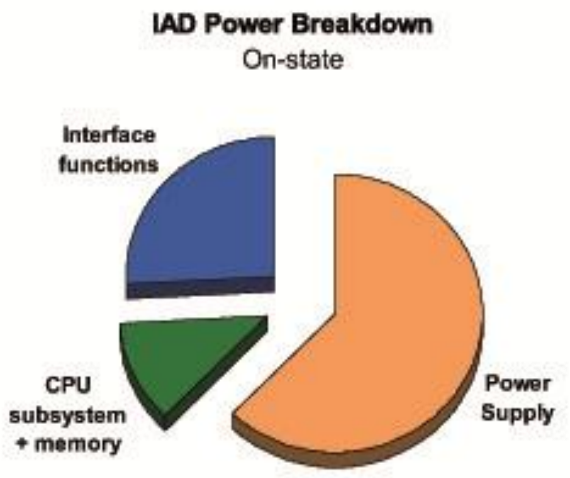

Figure 108: HG power consumption share for non-optimized power system, CPU subsystem and interface functions.

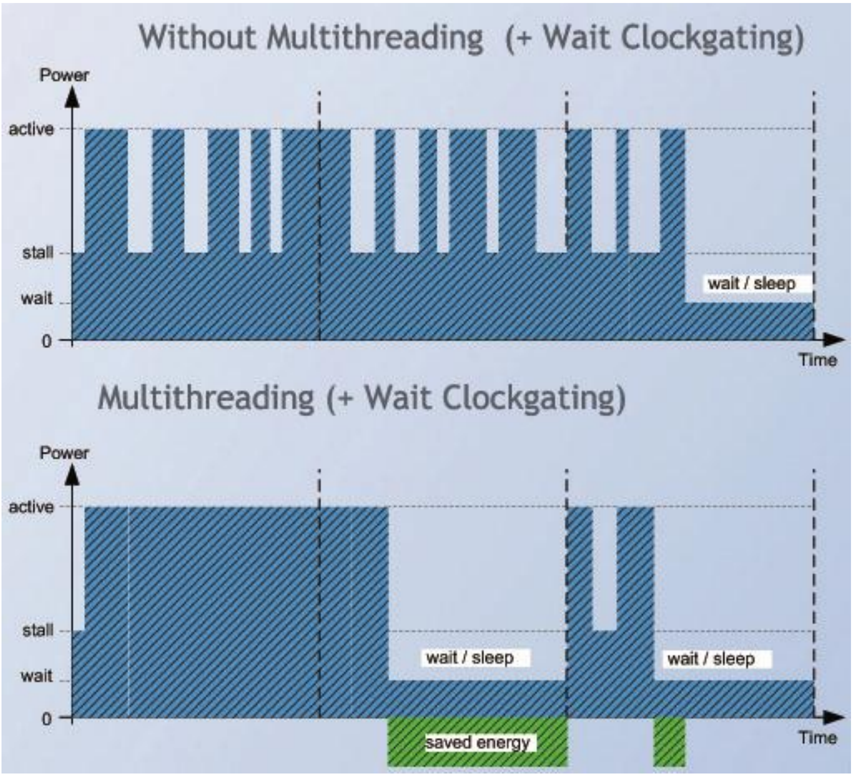

Figure 109: Multithread increases energy power saving in the processor.

The third part that can be optimized from the point of view of power consumption in the HG system is the processor and memory. Unfortunately, these are almost always active parts in the HG, since every operation rely on them. However, hardware techniques exist to optimize processor and memory. In particular, the multithreading technique proves to increase power saving since it permits to condensate processor jobs in the timeline and to have timeslots of rest, where the processor can go to wait/sleep states, increasing energy saving. Figure 109 shows how multithreading acts to save power. “Wait clockgating” is naturally essential so clocks can be partially disabled when sleep modes are in use.

It is important to cluster the jobs in timeslots, by the use of multi-threading, in order to have the possibility of obtaining bigger wait/sleep windows where the chip can go to low power states. In this way, the operations of saving context (e.g. CPU registers) are less executed, in particular only at the start and at the end of the rest windows, to restore context, thus reducing the overheads associated to the low power procedures.

Regarding the proposed targets for Home Gateways that is possible to reach within the ECONET project, in Table 35, the power consumption of a “triple-play” system in VDSL 17a, 802.11n 3x3, FXS, DECT, 4 Fast Ethernet ports configuration, is presented versus the EU CoC for HGs.

Table 35: Power consumption targets @ 230 V.

	Low Power State	On State
Home Gateway targets	6 W	11 W
CoC 11/12	6.52 W	13.45 W

Similarly, in Table 36, the power consumption of a “triple-play” system in VDSL 17a, 2,4 + 5GHz 802.11n 3x3, 2FXS/FXO, 2xUSB, DECT, 4 Gigabit Ethernet ports configuration, is presented versus the EU CoC for HGs.

Table 36: Power consumption targets @ 230 V.

	Low Power State	On State
Home Gateway targets	7 W	15 W
CoC 11/12	9.10 W	19.2 W

The next set of power values are the proposed targets for MSAN. They are smaller than those of the CoC, because of the use of vectoring. The set of power consumption targets is shown in Table 37 for each VDSL profile.

Table 37: Power consumption targets for each profile.

	Power consumption targets
VDSL2 profile 30a	1215 mW
VDSL2 profile 17a	895 mW
VDSL2 profile 8a	880 mW
VDSL2 profile 8b	1025 mW
VDSL2 profile 8c	710 mW
VDSL2 profile 8d	760 mW
ADSL2+	740 mW

5.3.3 Energy consumption profiles of FPGA based systems

The following section describes the power consumption estimation consumed on an FPGA system implementing Carrier Ethernet Switching based on three different scenarios to demonstrate the result of frequency power scaling. To examine the power scaling we took a Carrier Ethernet Switch implementation on Altera's Cyclone 3 80 FPGA with 24 Fast Ethernet ports and 2 GbE ports, where the packet processing and traffic management engines operate at a certain internal system clock, which is used as the scaling element. In the first scenario the system clock used was 100MHz and in the second scenario we decreased the system clock to 20MHz. Furthermore, to complete the picture we disabled the 24 fast Ethernet ports, such that the firmware on the FPGA employed only 2 GbE ports.

24G+2GbE with system clock of 100MHz

This configuration runs the Switch design in the FPGA with 24 Fast Ethernet ports and system clock of 100MHz. The main design block is the ENET core which embeds the network processing and traffic management functions, consumes 50,000 Logic Elements (#LUT) and runs at 100MHz, which is the system clock. Other entities are the MII, DDR controller, CPU/Host interface, together with other data interface including 2 x RGMII used for the 2 GbE ports and 3 x Octal SSSMIII used for the 24 Fast Ethernet ports.

For the purpose of this examination we will only play on the main ENET core that runs with the 100 MHz system clock, as the ENET core block consumes most of the power. Figure 110 describes the total power consumption estimation for the above configuration. It describes the total power consumption as a result of internal logic, internal memory (RAM), DSP (if in use), I/O and PLL.

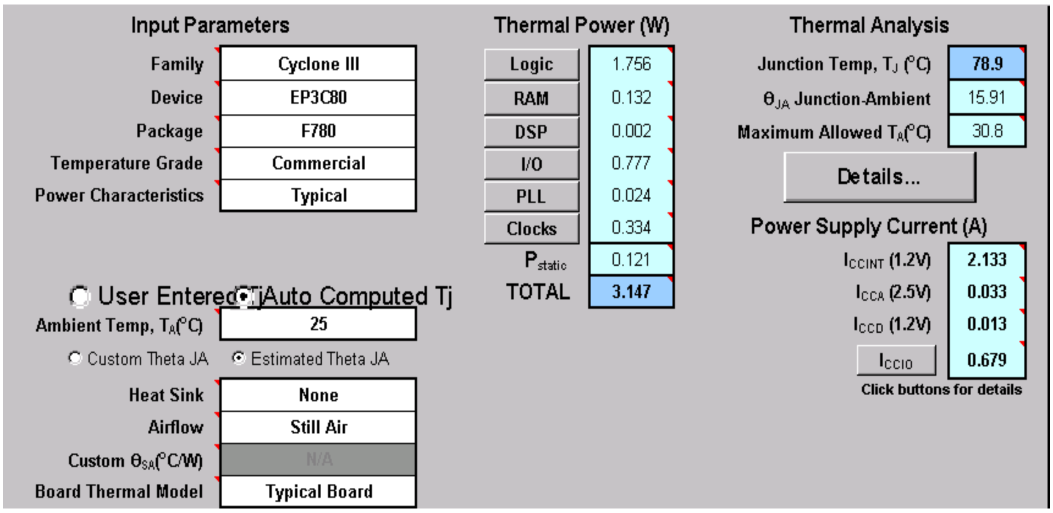

Figure 110: Total power consumption estimation for the ENET core with system clock of 100 MHz.

24G+2GbE with system clock of 20MHz

This configuration runs the Switch design in the FPGA with 24 Fast Ethernet ports and system clock of 20MHz. Compared to the above mentioned configuration, we only changed the system clock related to the ENET core from 100MHz to 20Mhz, and as a result of this change we see 80% power reduction on this element from 1.19W to 0.238 W, whereas the power consumption on the other elements remain the same. Figure 111 describes the total power consumption estimation for the ENET core in the 20MHz configuration. It demonstrates the power reduction in the logic due to the change in the ENET core system clock from 100MHz to 20MHz and reduction on the clock domain entity as a result of the change in the system clock domain frequency. We can also see that the consumption in other entities is the same as we did not change the actual logic or reduced the amount of memory, or I/O.

Figure 112 describes the total power consumption estimation for only 2GbE ports switch design that runs at 20MHz. In this configuration we eliminated the 24 Fast Ethernet ports. As a result, the design consumed less internal memory and the outcome is a power reduction of 20% on the RAM together with additional reduction in the logic and the I/Os. This configuration results in more than 50% reduction from the full functional 5Gbps Carrier Ethernet Switch that runs at 100MHz.

Based on the collected data for power scaling on FPGA based systems, Table 41 summarizes the total power consumption on the Altera's Cylcone 3 80 FPGA as a result of a system clock frequency scaling:

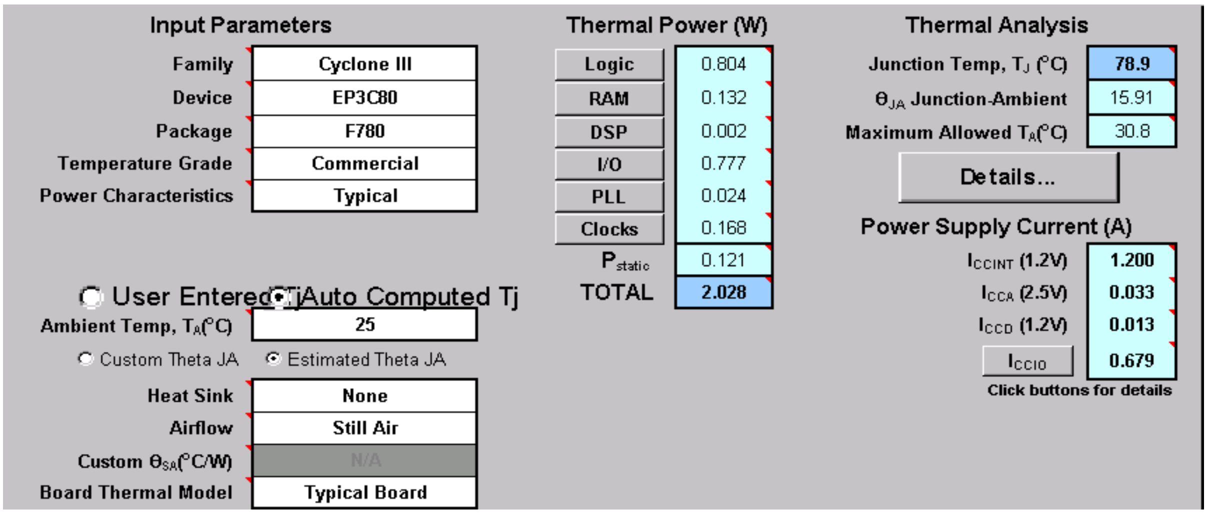

Figure 111: Total power consumption estimation for the ENET core with system clock of 20 MHz.

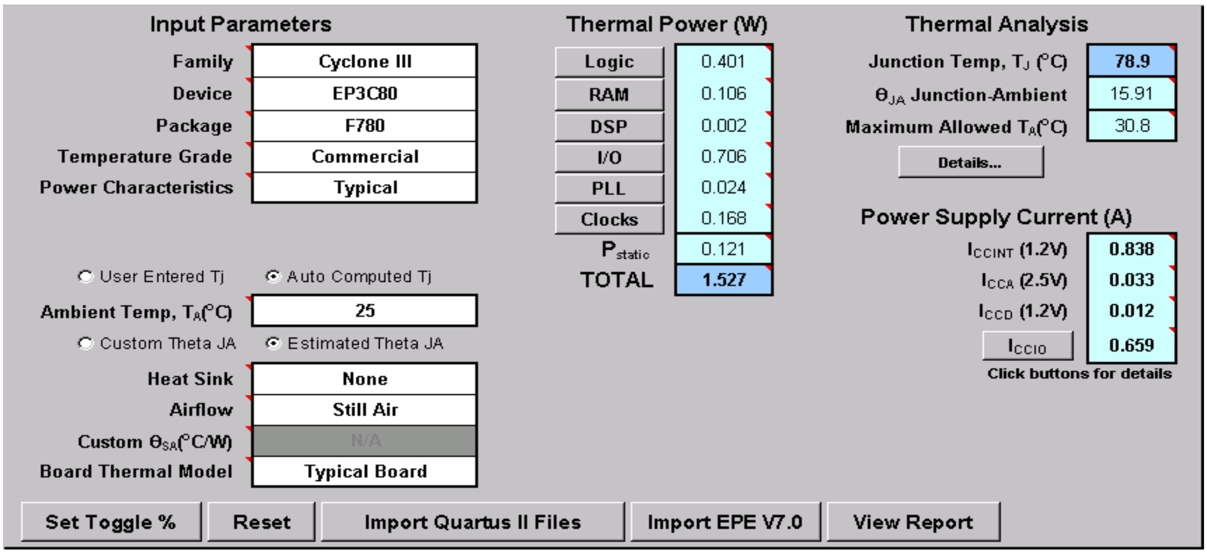

Figure 112: Total power consumption estimation for the ENET core with system clock of 20MHz for only 2GbE ports.

Table 38: Total power consumption on the Altera's Cylcone 3 80 FPGA as a result of a system clock frequency scaling.

System Clock	Power Consumption
100MHz	3.147W
75MHz	2.79W
50MHz	2.44W
33MHz	2.2W
20MHz	2.08W

5.3.4 Energy consumption profiles of software routers

As shown in section 4.1, silicon circuits at the data-plane are the most energy-harvesting parts in the largest part of network devices. Starting from this consideration, our first objective is to preliminary evaluate the impact and the feasibility of adopting power scaling mechanisms on packet processing engines, which represent one of the most energy-harvesting elements in many network platforms, in order to dynamically modulate energy absorption with respect to incoming traffic loads and service requirements.

To this purpose, we focus on promising new generation Software Routers (SR) [60,61], based on open source SW and on multi-core Commercial Off-The-Shelves (COTS) hardware since, unlikely the largest part of commercial platforms:

- Their HW already provides advanced power management capabilities (PSM) [62],[46] by means of the Advanced Configuration and Power Interface (ACPI) technology [63];
- Their SW is open and can be modified for enabling new functionalities/mechanisms.

In this scenario, we pursued a "bottom-up" approach: at first, we try to enable energy management capabilities in the SR data plane, especially in terms of frequency scaling support at packet processing engines. Then, our objective is to deeply understand their impact by characterizing the trade-off between SR performance and energy requirements under different system configurations. With this main aim, we performed several benchmarking sessions with a heterogeneous set of HW platforms. These measures gave us the chance both to understand the relationship between the SR internal dynamics and its external/power-related performance, and to derive a simple model to estimate SR behaviour with respect to different setups and traffic offered loads.

Regarding power scaling HW capabilities, we decided to evaluate the performance provided by a heterogeneous set of PC platforms, based on different CPU models and ranging from high-end servers to common workstations and mobile architectures (see Table 39). For each selected platform, we performed two types of power consumption measurements:

- Φ_{idle}- Idle measures: in order to evaluate the average power consumption when all the Cores in the system are idle, and do not perform any tasks;
- Φ_{active}- P_n state measures: to quantify the average power consumption when one or more Cores are running (i.e., performing SW operations for 100% of their time) at a fixed P_n state.

While the first kind of measurement is useful to evaluate the minimum power consumption for each selected platform, the Φ_{active} measures point out how much energy a Core can absorb when it is fully active at a certain ACPI P_n state. In this respect, these two measures can be thought of, respectively, as the upper and lower bounds of SR performance when deployed in a real network.

We also decided to consider an additional performance index, useful to evaluate the Core's computational capacity at the given frequency. Since data-plane functionalities are certainly the most CPU intensive and time consuming operations to be carried out in a SR, we decided to consider the maximum throughput, in terms of forwarded packets per second, as the key performance index in order to evaluate CPU computational capacity.

Figures 113 and 114 show the Φ_{active} and Φ_{idle} results obtained with a heterogeneous set of HW architectures, which includes high-end servers, workstations and mobile architectures. In detail, we extensively tested eleven HW platforms with the characteristics reported in Table 2. Figures 113 and 114 highlight the gap, in terms of power consumption and forwarding performance, between the different HW platforms. These differences obviously depend not only on the specific CPU model, but also on the other HW components (e.g., mainboards, etc.). For the sake of simplicity, we drop this additional power consumption contribution in the following discussion, since it does not affect the obtained results and conclusions.

Table 39: Considered HW platforms for Linux SW Router development.

Name	Features
HW platform 1 (X)	dual CPU Intel XEON Dual Core 5050 "Dempsey", 3.0GHz, 64bit, RAM 533 MHz, PCIe 16x
HW platform 2 (PD)	Intel Pentium-D 925Dual Core, 3.0GHz, RAM 667 MHz, bus PCIe 8x
HW platform 3 (A)	AMD Athlon 64 dual Core X2 3800+, 2.0GHz, RAM 667MHz, PCIe 16x
HW platform 4 (PM)	Intel Pentium-M 745 "Dothan", single Core 1.8GHz, RAM 266MHz, PCIe 8x
HW platform 5 (O)	dual CPU AMD Opteron Dual Core 265, 1.8GHz, RAM 400MHz, PCIe 16x
HW platform 6 (C2)	Intel Core 2 Duo E2160, 1.8 GHz, RAM 667 MHz, PCIe 8x
HW platform 7 (A2)	AMD Athlon 64 3500+, 2.2 GHz, RAM 200 MHz, PCIe 16x
HW platform 8 (X2)	dual CPU Xeon Quad Core X5550 "Gainestown", 2.66 GHz, 64 bit, RAM 1333 MHz, PCIe 16x
HW platform 9 (X3)	dual CPU Xeon Six Core X5650 "Gulftown", 2.66 GHz, 64 bit, RAM 1333 MHz, PCIe 16x
HW platform 10 (I)	Intel I5 Quad Core 750 "Lynnfield", 2.67 GHz, 64 bit, RAM 1333 MHz, PCI 16x
HW platform 11 (IA)	Intel Atom Dual Core D510 "Pineview", 1.83GHz, 64 bit, RAM 800 MHz, PCI 16x

High-end server architectures (i.e., X, X1, X2, and O) obviously show both a heavy power consumption (more than 150 Wh), and a high forwarding performance, especially when a large number of Cores is employed. On the contrary, workstation architectures (PD, A, I and A2) exhibit an intermediate performance level, but both power consumption and forwarding performance can heavily change, according to the selected HW platform. For example, the PD platform is characterized by a much higher power consumption than the A and the A2 ones, while the A platform, thanks to its high-speed internal busses, shows remarkable values of packet throughput (about 1.5 Mpkt/s).

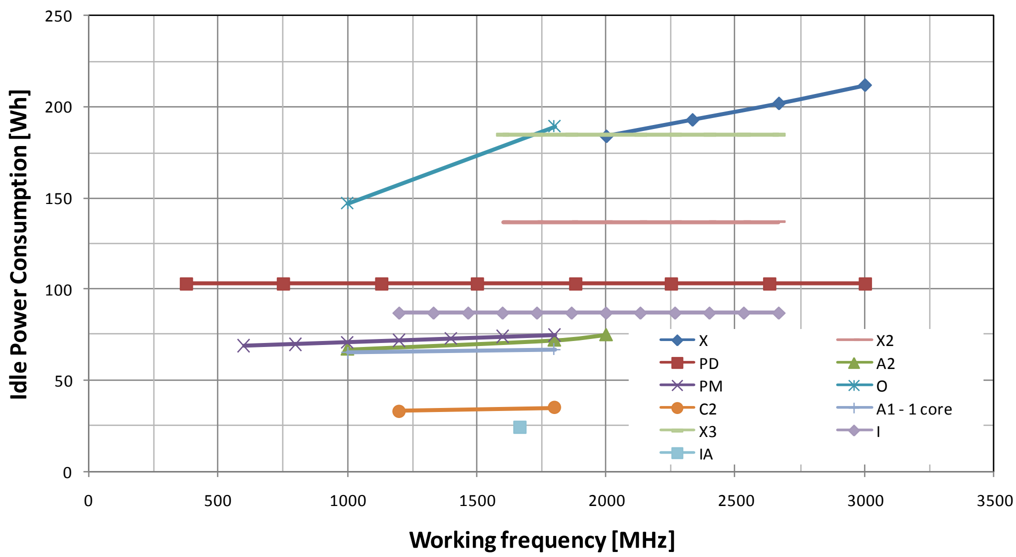

Figure 113: Idle power consumption for all the selected HW platforms.

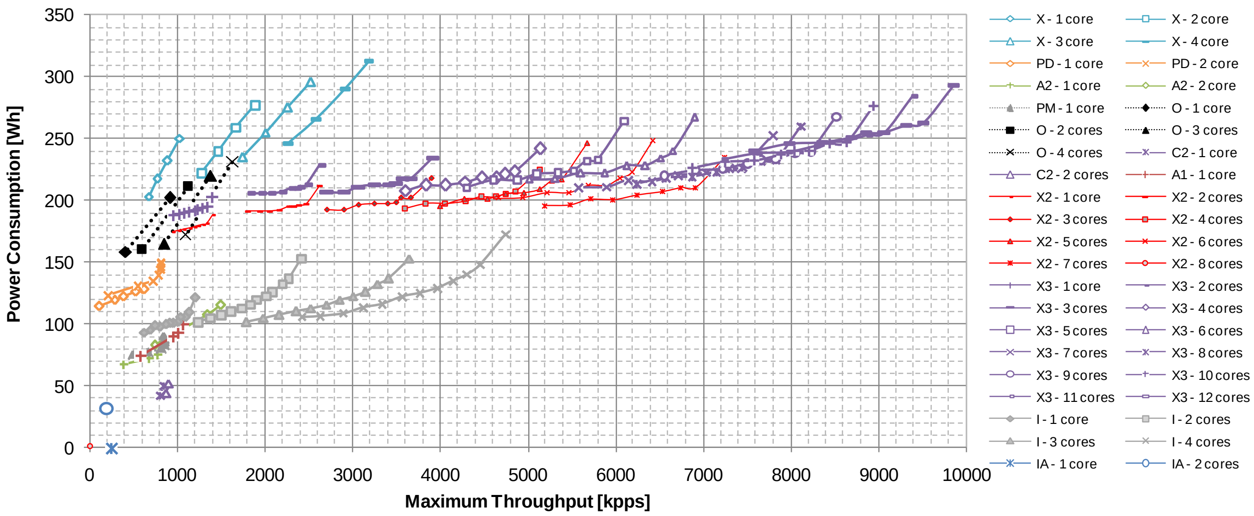

Figure 114: Power consumption and maximum throughput values for all the selected HW platforms. For each HW platform, the tests were performed by changing the Core *P* state, and with a variable number of active CPUs/Cores.

In particular, the forwarding performance of the PD and the A2 platforms are limited by the bandwidth and the efficiency of internal busses, which do not allow exploiting the full computational capacity of CPUs/Cores.

Similar remarks can be made also for the mobile architectures (i.e., PM, IA and C2). In detail, we can note how the C2 HW platform provides a very low power consumption, which ranges between about 30 and 50 Wh.

Besides interesting information on different HW platforms' behaviour, Figures 113 and 114 show that all the analysed PSMs (both in AMD and Intel processors) allow to linearly scale the maximum forwarding rates with respect to power consumption. This linear behaviour is maintained not only for different equivalent frequency values, when a single core is used for the SR data plane, but also for a rising number of CPUs/Cores.

5.4 Monitoring and management issues

In this subsection, an initial set of energy aware parameters is defined and specified in order to be exposed to control and OAM applications for the autonomic and/or user driven configuration and monitoring. Monitoring and management of energy aware parameters represents a key point for the ECONET project, since it aims at introducing and developing novel integrated frameworks for the energy-aware network control, management and monitoring, working on top of the devices green capabilities and the abstraction layer interface.

Monitoring may be related with the local status of the devices or the network status and may be applied in a centralized or distributed manner. The monitoring process consists of gathering metrics about the state and the power consumption of the hardware resources available on each network device which is part of the monitored infrastructure. Two reference models are available: the push and the pull model.

- The push model is polling based: the monitoring process samples the energy parameters periodically, acting explicit monitoring request to the network node.
- In the pull model, each network node sends periodically new values of the consumed parameters to the monitoring system.

The pull model is performed by additional monitoring agents installed on each node, which will afterwards send the gathered data to a collector for processing and persistence. In case of push monitoring, multiple frameworks can be used, such as Ganglia[12] (active monitoring) and CoMo[13] (passive monitoring). These frameworks differ in the communication protocol employed in transferring gathered metrics, the architecture of the monitoring strategy subsystems, the level of monitoring resources and the exposed API for accessing the collected information.

In case of physical hosts, information about the energy consumption could be obtained in different ways, depending on the available hardware capabilities. There are four possible sources for such information:

- Static datasheet information, when there is no possibility of measuring or estimating the energy consumption.
- Estimation of the power consumption using the data monitored, like utilization values for network elements, boards and network links, through the use of a linear power model. Such a model would use the idle power value and increase the value proportionately with the usage factor of the aforementioned resources. This information could be calculated on each physical host and then injected in the monitoring system (e.g. Ganglia) for retrieval by the Green Abstraction Layer.
- Direct measurements from host's energy aware hardware.
- Direct measurements of instantaneous power consumption from an attached power monitor. Such data collected from the power monitor could be fed to an energy estimation framework like AMEE, allowing calculation of different statistics or it could be processed directly on the host.

For implementation, one of the first three mechanisms could be selected, as the fourth implies the use of a power metering system inside each device (wattmeter), which is impractical for all

[12] http://ganglia.sourceforge.net/
[13] http://como.sourceforge.net/

ECONET physical hosts. However, for the estimation model to function properly, parameters of the model would need to be determined using some sort of external power monitor.

Energy aware parameters

The energy efficiency support gathers multiple metrics, like: current (instantaneous) power consumption, average / max / min power consumption over last minute / hour / day / month or, optionally, over a custom period (which implies storing the power usage data). This data would then be made available on demand (through adding of custom probes) or, it could be fed into the monitoring system as current power consumption. Table 40 contains the first proposal of parameters that should be monitored, some of which were proposed within RFC4258.

Table 40: Monitored parameters.

Parameter name	Description
ResourceId	Identifier of the network resource
Status	Represents the power status of the network device. Possible values are [PS0=SwitchOn][PS1][…][PSN=Switch Off]
TransitionTime	The average time necessary to the specified device to switch between two different states.
Mode	Represents the operational mode of the power status, related to the specified network device. Possible values are [static][dynamic] - In Static mode, the user is able to configure the power status of the managed devices, manually (switch on, switch off, etc ..) using management tools. - In Dynamic mode the power status of the selected interface is managed automatically, without any configuration steps. In this scenario, an appropriate algorithm analyses the power consumption of each device of the network element and calculates the best path minimizing the costs associated to it. Basic algorithms, interfaced with control planes facilities, will derive from PID, Sliding Mode, Optimal Control and Robust Control standard algorithms.
PowerInfo	Average power consumption of the selected network device
PowerSource	PowerSource selects the source of the power consumption data, related to a well-defined network device. - In metering mode, the average power consumption of each interface/device is calculated using a real-time metering system connected directly to the selected device. - In database mode, the average power consumption of each interface/device is retrieved from a local database, containing the power info of each device, statically.
Transition from sleeping status	The time needed for the device to transit from a sleeping to an active status. This parameter is crucial for the design of traffic engineering and distributed optimization techniques.
Transition time for rate adaptation	According to the network conditions, rate adaptation may be applied for reduced energy consumption. The transition time among different rates is important for the design of the rate adaptation strategy. Furthermore, in case that the temperature of the chipset is too high, the data rate can be reconfigured to a lower one in order to reduce the power consumption and temperature of the device.
Power efficacy	The percentage of device mechanisms that are currently working in power saving mode. For example, on a modular switch it can be the percentage of modules currently working in power efficient mode.
Power capabilities	A well-defined list of power capabilities that the device supports. For example EEE, interface standby, system standby.
Chipset temperature	Monitoring the temperature of the chip-set. Once the temperature is too high, the feature can be reconfigured or some channels can be disabled or the datarate can be reconfigured to a lower one in order to reduce the power consumption and temperature of the device.
Traffic counters	Embedded traffic level counters useful to keep count of traffic.

5.5 Summary

Section 5 detailed the technology specifications of energy efficient devices. A detailed description of the energy aware design space has been provided for future energy aware technologies at the network device data plane and at the network control plane. At the network device data plane the techniques include total re-engineering in order to cut down on power, dynamic adaptation of operating rate with response to current load, various sleeping techniques for devices that can go offline when the current load can be handled by a subset of them and energy efficient Ethernet, a new key technology appearing at various parts of the network. At the control plane, additional techniques include energy aware traffic engineering for selecting a subset of paths to carry the traffic, various joint resource (re)allocations around energy consumption and performance and virtualization techniques for adapting the function of network elements on the fly.

Upon the description of the future energy-aware technologies, the power benefit expected from applying these technologies in network devices in the different network parts have been described. Specific consumption profiles have been presented for switch systems, VDSL home gateways, FPGA based systems and software routers.

In switch systems, it was shown that by applying smart adaptation of power consumption of each port to the actual traffic the switch serves, we can accomplish power reduction up to 50% of the maximum power consumed by today's switching systems. By applying smart port shut-down techniques, we can achieve power reduction of up to 50%, while by applying smart cooling techniques we can reduce the average working point of the fans to 50% of max speed, and by doing so achieve additional 10% reduction in total power consumption. Additional significant power improvement can be achieved through the integration of power management and control tools that will have additional impacts on the total power consumption of the whole system.

For Home Gateways (HGs), it is possible to achieve power saving turning to standby modes unused interface functions during the fruition of one particular HG service. The power supply system is identified as the more power hungry part of the HG, thus it has to be also optimized. By the use of optimized power designs it is possible to achieve efficiency of 80% for the external power supply and of 80% for the internal DC-DC converter, thus having a total 40% reduction in power losses compared to old designs. The third part that can be optimized from the point of view of power consumption in the HG system is the processor and memory. For this purpose, the multithreading technique proves to increase power saving since it permits to condensate processor jobs in the timeline and to have timeslots of rest, where the processor can go to wait/sleep states, increasing energy saving.

The power consumption estimation on an FPGA system implementing Carrier Ethernet Switching was also presented based on three different scenarios to demonstrate the result of frequency power scaling. Based on specific configurations, a minimum power reduction of 20% on the RAM together with additional reduction in the logic and the I/Os is achieved based on a switch design that runs at 20MHz. This configuration results in more than 50% reduction from the full functional 5Gbps Carrier Ethernet Switch that runs at 100MHz.

In network routers, we focused on promising new generation Software Routers (SR), based on open source SW and on multi-core Commercial Off-The-Shelves (COTS) hardware. We tried to enable energy management capabilities in the SR data plane and understand their impact by characterizing the trade-off between SR performance and energy requirements under different system configurations. For this purpose, several benchmarking sessions were performed with a heterogeneous set of HW platforms. These measures gave us the chance both to understand the relationship between the SR internal dynamics and its external/power-related performance, and to derive a simple model to estimate SR behaviour with respect to different setups and traffic offered loads.

Finally, monitoring and management issues were addressed, where an initial set of energy aware parameters is defined and specified in order to be exposed to control and OAM applications for the autonomic and/or user driven configuration and monitoring.

6 Target energy saving and general performance indexes

The introduction of energy saving mechanisms involves the need of commonly accepted criteria to evaluate their effectiveness, thus the standardization of a set of benchmarking methodologies and performance indexes. This section initially provides a brief state of the art of the existing standards, in order to point out their strongest and weakest points. Then, the evaluation methodology that will be followed within ECONET is proposed, describing both the performance indexes conceived to denote the trade-off between power consumption and network performance and the set of test procedures used to obtain the desired indexes. The most innovative features introduced are the use of latency to characterize the performance degradation due to the power management methods, the use of bursty traffic in order to perform more realistic tests and thus stay closer to a real Internet scenario and the examination of the trade-off among the energy consumption in all the network devices and the QoS perceived by the end-users. Evaluation criteria are defined for both data and control plane, focusing mainly on the device and its components or the entire network accordingly. It is important to note that in the evaluation methodology we focus on the energy consumption of devices that constitute the network without looking in detail in the energy consumption of the peripheral devices, like air conditioning or power supplies, whose energy models are generally rough, if not unknown altogether.

In addition to the evaluation criteria, an analytical framework is introduced for representing the impact of energy adaptive technologies and solutions for network devices under investigation by the ECONET consortium. The framework is applied in the ECONET reference scenarios and the energy saving gains that can be achieved based on the application of energy-aware techniques are estimated.

6.1 Performance indexes and benchmarking methodologies

This section is meant to discuss and compare the three standard methods for measuring power consumption at varying performance levels of network devices and systems. In the following analysis, each method will be investigated singularly, in order to establish in a complete and, as much as possible, easy way, its specific behaviour. Further considerations and comparisons will be provided at the end of the section. Although the last method here described has a different aim, as it does not perform any power consumption measure, it is included because it represents a valid guideline for both performance testing and set up procedures. Moreover, some of the tests it defines are already exploited by the other standards to select the starting parameters for further measurements.

6.1.1 Existing Standards and Comparison

Energy Consumption Rate (ECR 3.0.1 Standard)

ECR 3.0.1 [26] is a proprietary standard owned by Juniper and IXIA. It allows measuring the ratio between power and performance as well as the consumption due to single components (in presence of redundancy) and the system energy consumption over a projected lifetime. The main test, defined as mandatory, is performed measuring the maximum offered load the device under test can support without losing any data. A router tester is used for this purpose. Traffic is sent at that speed for 1200 seconds, and a watt-meter measures the average power consumed in this time interval. Traffic is then sent at 50%, 30%, 10% and 0% of the maximum throughput, and average power consumption is measured during each transmission. The chosen packet size is the maximum supported by that kind of the System Under Test (SUT). The Energy Consumption Rating (ECR) index can be expressed as follows:

$$ECR = \frac{\Phi_{100}}{T_{100}} \quad (5)$$

where Φ_{100} represents the average power consumption at maximum throughput and T_{100} is the throughput itself.

The ECR index has the clear and simple aim of quantifying how much energy is needed to forward one Gigabit of data at the maximum speed. This performance index is clearly very “unnatural”, since it evaluates the SUT at an operating condition (maximum speed with Constant Bit Rate (CBR) traffic flows) that is seldom (or never) reached in real-world deployment scenarios.

Notwithstanding these considerations, the ECR index can be obviously useful to evaluate the SUT “green” re-engineering level. Unfortunately, the ECR does not include any indications regarding power management mechanisms possibly included in the SUT, since their effects become evident only at low levels of traffic offered load. However, the standard introduces an optional test aimed at estimating a version of the ECR index, called ECR-VL (ECR-Variable Load), weighted for various offered load (0%, 30%, 50%, 100% of the effective throughput). All the benchmarking tests use CBR traffic flows with fixed-size packets.

Telecommunications Energy Efficiency Ratio (ATIS060015.02.2009 Standard)

As far as the testing methodologies are concerned, the ATIS 060015.02.2009 standard looks very similar to the ECR 3.0.1. However, it is addressed to transport/core network products, owned by the carrier and explicitly characterized by redundancy. The evaluation metric that is introduced is an index called Telecommunication Energy Efficiency Ratio (TEER) that can be obtained by putting together the results of each component, collected in a database. It is also possible to provide a series of certified configurations, each one characterized by its index.

For each component, the database maintains the maximum throughput and the average power consumed at 100%, 50% and 0% of that throughput. TEER is obtained as the ratio between the sum of the maximum rate for each component and the sum of the consumptions, for each component, at the three traffic loads previously mentioned. In order to simulate real world behaviour, Internet MIX (IMIX) traffic profile is used for the transmissions, which means that packets of different lengths are generated by the test equipment according to a statistic distribution and transmitted to the SUT.

The main difference between the two standards can be seen in the amount of produced results: ATIS 060015.02.2009 only provides TEER, while ECR 3.0.1, aside from the index obtained with the mandatory test, gives a wider data volume, including average consumption measured on different traffic rates. The devices taken into consideration by the second standard, however, are more specific, and they focus mainly on modularity.

Benchmarking methodology for network interconnect devices (RFC 2544)

Notwithstanding the RFC 2544 is not oriented to evaluate energy-efficiency, it is included it in this section, since it is the "mother" of all standards for benchmarking network devices. Thus, recalling its base concepts may be useful to extend ECR 3.0.1 and ATIS 060015.02.2009.

This standard provides a very complete set of methodologies and performance indexes for evaluating network performance. The same tests can be applied to all kinds of network interconnection devices, in all their possible configurations. The following tests can be taken into consideration:

- *Throughput Test*: the throughput is the fastest rate at which the count of test frames transmitted by the SUT is equal to the number of test frames sent to it by the tester. For each available frame length the maximum throughput is calculated; this provides not only the first test result, but also a binding datum for the next measures.
- *Latency Test*: the procedure that determines latency uses the previous result as a starting point. Latency represents the interval between the time at which a frame is fully transmitted and the time at which the same frame is received.
- *Back-to-back Test*: bursty traffic characterized by the maximum burst length allowed is generated by the tester. The back-to-back value is the number of frames in the longest burst that the SUT can handle without losing any data.

Comparison and extensions

The main features and performance indexes obtained with the considered standards are summarized in Table 41.

Table 41: Comparison of the Analysed Benchmarking Standards.

Standards	Benchmarking scenario	Main results obtained
ECR 3.0.1	1. CBR traffic 2. Traffic at different loads 3. Maximum packet size	1. ECR [W/Gbps] 2. ECR-VL [W/Gbps], ECR-EX [W/Gbps] 3. Energy Bill Estimates [$]
ATIS 060015.02.2009	1. CBR traffic 2. Traffic at different loads 3. IMIX traffic	1. DTEER [Mbps/W] 2. CTEER [Mbps/W]
RFC 2544	1. CBR traffic 2. Bursty traffic 3. Traffic at different loads 4. Range of packet sizes	1. Maximum throughput [Gbps] 2. Max, average, min latency [us] 3. Maximum burst length [bursts/frames]

Based on the above-mentioned standards, it can be concluded that the existing standards present some interesting features, but at the same time they lack important indications that would be necessary for a thorough evaluation of future network devices. Considering ECR 3.0.1 and ATIS 060015.02.2009, both of them are focused only on the measurements of energy consumption, however, they lack clear indications on evaluating the trade-off between power absorption and network performance. In addition, since none of their methodologies takes into consideration bursty traffic, the obtained results cannot provide an accurate representation of real traffic. Regarding the RFC 2544, the main characteristics regard its multi-purpose nature: although it is not designed for evaluating energy efficiency, its concepts and methodologies can be deeply exploited within other contexts. The fact that it can be applied to all kinds of network systems, embracing different protocols, topologies and traffic nodes testifies for its versatility.

Furthermore, latency is not taken into consideration neither in ECR 3.0.1 nor in ATIS 060015.02.2009, as they both use only throughput to characterize performance. However, in energy-aware systems power reduction tends to cause an increasing of service times, generating delays. For this reason, a correct exploitation of energy efficiency techniques cannot be evaluated regardless of latency. In addition, both standards employ CBR traffic to determine the average power consumed during the transmission. This assumption can be unrealistic if we consider the fact that real traffic generally has a bursty composition. However, CBR traffic could be considered as a limit case and included in a wider set of measurements. Regarding results' representation, determining a global index would be desirable, since it is more synthetic and gives an immediate idea on a system's behaviour. On the other hand, some critical aspects could be neglected without any further information.

Starting from these considerations, the methodology that will be followed within ECONET consists in determining two indexes: one represents the amount of energy that can be saved through power management capabilities, while the other one quantifies the performance degradation that these capabilities can provoke. These indexes will be explained in a more detailed way in the next subsections and will be correlated with specific parameters related with the data and the control plane.

6.1.2 Evaluation of energy-aware data plane

The evaluation of the energy aware data plane has to take into account the improvement in energy efficiency within each device that is achieved through the application of energy-aware techniques in this plane (as they are already detailed in section 5.1). In order to be able to proceed to evaluation, specific performance indexes have to be defined taking into account the existing ones' in the available standards. For instance, the index provided by ECR 3.0.1 shows in a simple way the amount of energy consumed to forward one Gigabit of data. For this reason, it can be usefully taken into consideration to represent how a device reacts to a limited behaviour of traffic, as in the case of CBR flows.

The information obtained with this result becomes not sufficient especially in the presence of power management policies. Both adaptive rate (AR) and Low Power Idle (LPI) techniques have the drawback of increasing latency. Power scaling causes a stretching of packet service times, while idle logic introduces an additional delay in packet service, due to wake-up times. Moreover, power scaling causes larger packet service times, and consequently shorter idle periods. For all these reasons, it is necessary to collect and put data together to quantify what we gain in terms of energy saving and what we lose in terms of service delays.

Performance indexes

Energy gain represents the power saving obtained thanks to power management in comparison to a scenario with no such capabilities:

$$\Phi_{\%} = \frac{\Phi_{max} - \Phi_c}{\Phi_{max}} \quad (6)$$

where Φ_c is the current power consumption and Φ_{max} is the maximum consumption reached by the device. The result gets closer to 1 as consumption decreases.

Performance degradation is expressed as a ratio between the values of packet latency in the case of an ideal network device (i.e., with an infinite processing capacity), and the ones measured with the real SUT:

$$L_{\%} = \frac{\tilde{L}_i + \frac{1}{2}\left(L_i^{max} - L_i^{min}\right)}{\tilde{L}_r + \frac{1}{2}\left(L_r^{max} - L_r^{min}\right)} \quad (7)$$

where parameters with the *i* index represent the latencies for the ideal device, and the ones with the *r* index are measured on the real SUT. $L^{min}, \tilde{L}$ and L^{max} are the minimum, average and maximum values, respectively, of packet latencies. The values for minimum, average and maximum latency in the ideal case have to be computed by starting from the packet transmission times on input and output links. It is worth noting that if the packet transmission time on the input link is longer or equal to the one of the output link, then $L_i^{min} = \tilde{L}_i = L_i^{max}$. Otherwise, as shown in Figure 115, these latency values have to be estimated by considering traffic burstiness in an explicit way.

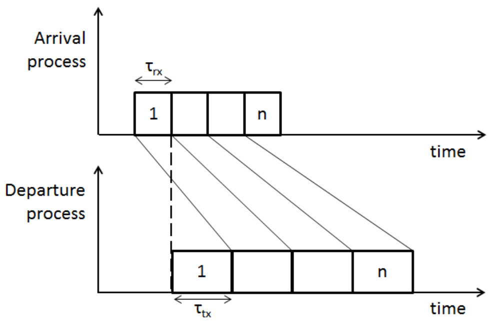

Figure 115. Packet burst processing in an "ideal" SUT

We can simply compute the L_i^{min} value as the sum of the packet transmission times on the input and output links (τ_{rx}, τ_{tx}):

$$L_i^{min} = \tau_{rx} + \tau_{tx} \quad (8)$$

With reference again to the simple example in Figure 115, the L_i^{max} parameter can be calculated as follows:

$$L_i^{max} = \tau_{rx} + n\tau_{tx} - (n-1)\tau_{rx} = n\tau_{tx} - (n-2)\tau_{rx} \quad (9)$$

where *n* is the length of the incoming bursts. Thus, the $\tilde{L}_i$ value can be simply expressed as:

$$\tilde{L}_i = \frac{n+1}{2}\tau_{tx} - \frac{n-3}{2}\tau_{rx} \quad (10)$$

Benchmarking methodologies

As in the in ECR and ATIS standards, it is suggested to start benchmarking tests by exploiting the RFC 2544 Throughput Test. But differently from the above methodologies, we propose to perform test sessions with:

- different loads, namely 100%, 50%, 30%, 10% and 0% of the output link bandwidth (and not of the effective throughput – this will make more comparable results obtained with different SUTs or with different configuration of the same SUT);
- different packet burst sizes, namely 1 (CBR flows), 10, 50, 100, 200 packets;
- different packet sizes (as suggested by the RFC 2544), and where possible with IMIX profiles.

It is worth noting that, when possible, packets shall have different IP source and destination addresses, since devices often include multiplexer and de-multiplexer elements working on a per-flow basis. In this way different IP couples can be classified as different flows, and then trigger different parallel HW pipelines internally to the SUT. On each session, the average power has to be measured. Regarding network performance different values can be collected: the average throughput, the minimum, average and maximum value of packet latency, and packet loss rates. If the SUT has the possibility of entering various power management configurations, measures are repeated for each of them.

6.1.3 Evaluation of energy-aware control plane

The steps towards improving the energy efficiency of an already implemented network involves replacing old equipment with new energy friendly one and/or performing network consolidation at each PoP reducing the number of energy consuming devices and/or changing network topology and/or developing control strategies that rely on the overprovisioned/redundant nature of networks to reroute traffic while putting devices to sleep and/or developing strategies that exploit the power aware capabilities of the devices in expense of their performance.

The development of strategies at the control plane facilitates the deployment of power aware capabilities of the equipment as well as the overprovisioned and redundant nature of the implemented networks to improve the energy efficiency of the network. Strategies are targeted to achieve energy savings on a reference network reducing its total energy consumption over time while managing effectively serving the same traffic as in the energy agnostic and performance oriented operation without degrading the quality experienced by the end user.

Performance indexes

Let $P_{N_R}(t)$ be the total instant power consumption at time t over a reference network N_R, defined as the sum of the instant power consumption $P_{D_i}(t)$ at time t of all n involved in the network operation devices D_i where $i \in [1..n]$;

$$P_{N_R}(t) = \sum_{i=1}^{n} P_{D_i}(t) \; : D_i \in N_R, i \in [1..n]$$

Considering a time period T then the total energy consumption E_{N_R} over the reference network N_R for the period T is defined as

$$E_{N_R} = \int_{t=0}^{T} P_{N_R}(t) dt$$

Let $R_{N_R}(t)$ be the aggregate traffic rate at time t at the reference network N_R, defined as the sum of the traffic rates $R_{f_i}(t)$ of all m distinct end-to-end flows f_i , $i \in [1..m]$ at time t;

$$R_{N_R}(t) = \sum_{i=1}^{m} R_{f_i}(t)$$

Considering a time period T then the total volume V_{N_R} inserted into the reference network N_R for the period T is defined as

$$V_{N_R} = \int_{t=0}^{T} R_{N_R}(t)dt$$

Considering that power is expressed in *Watts*, volume in *bits* and time in *seconds*, network N_R requires E_{N_R} *Watts* for serving V_{N_R} *bits* of aggregate traffic volume over a time period of T *seconds*. Let $\bar{e}_{N_R}$ be the average energy in *Watts* required for serving 1 *bit* (average *energy-per-bit*) over the network N_R considering an operation period of T *seconds*;

$$\bar{e}_{N_R} = \frac{E_{N_R}}{V_{N_R}}$$

Let N_R^S be the network after applying energy aware strategies and $E_{N_R^S}$ the new total energy consumption over a period T for *the same traffic volume and pattern over time*. Let $\bar{e}_{N_R^S}$ be the new average energy-per-bit achieved over the period T. Let $P_{N_R^S}(t)$ be the total power consumption at time t over the network N_R^S. The energy saving ratio achieved $S_{N_R^S}$ is defined as

$$S_{N_R^S} = \frac{E_{N_R^S}}{E_{N_R}} = \frac{\bar{e}_{N_R^S}}{\bar{e}_{N_R}}$$

An objective comparison between the energy consumption on the reference network and the energy consumption on the network after applying energy aware strategies in fact would require a thorough comparison considering several workloads and then probably considering energy saving as a weighted average over all energy savings at the various workloads.

Traditional network performance metrics such as end to end delay, throughput and jitter that reflect offered QoS should be measured for both the reference network and the network illustrating the developed strategies and compared towards deciding performance degradation against achieved energy savings. However developed strategies may also include methods to conceal from the end user the effects of a probable QoS degradation. Thus, Perceptual QoS (PQoS) – or Quality of Experience (QoE) - as a metric of the quality experienced by the end user should be also introduced and evaluated against classic QoS and achieved energy savings; a high PQoS translates to a high performance experienced by the end user irrespectively from the fact that QoS metric may indicate otherwise. Most commonly PQoS is related to the type of service considered.

Furthermore, an indicator of how well developed strategies perform is the level of energy proportionality to the network load they reach. Although this indicator is to be further defined, considering as the network load the traffic injected into the network from several origins, it seems that considering the average energy consumption versus the network load as this is generated as the sum of flow rates of several random combinations of origin-destination pairs over a time period, would provide valuable info regarding the energy proportionality achieved.

Specific bounds can be defined for this purpose. For an appropriately chosen observation period and a network topology illustrated by a set of devices, the upper bound to energy consumption is given when the fully energy agnostic case is considered where operation is targeted for maximum performance (best offered QoS) and redundancy, irrespectively of the energy consumption. Regarding the theoretical lower bound to energy consumption this is given when:

1. the fully energy proportional case is considered where ideally it is considered that each device consumes energy directly proportionally to the traffic it serves; it consumes zero energy when serving no traffic while it reaches its maximum consumption when serving 100% of its maximum throughput, and
2. every origin-destination communication follows the most energy efficient path.

Developing strategies for a considered reference network aims at approaching as close as possible the lower bound to energy consumption - i.e. approaching the upper bound to energy savings-, while managing at the same time to retain the performance perceived by a user.

Benchmarking methodologies

Evaluation of the proposed strategies can be conducted based on simulations and on implementations in the ECONET testbeds. A fully simulated environment permits studying their behaviour in a broad space of diverse conditions and reach safe results while the next step is to evaluate them over prototypes before implementing them in a real working environment. In all the cases, profiling of the traffic over a reference network is crucial for finding patterns and behaviours that the strategies under development may take advantage of to save energy. Although evaluation of new strategies may only consider workloads and traffic patterns met in the reference network, evaluation over several workloads and traffic patterns synthetically generated would be necessary to exhibit the generality of the proposed strategies and reveal the relation between performance and area of deployment in terms of workload and traffic characteristics. Furthermore, the achieved energy saving should be examined versus QoS and PQoS while the achieved energy proportionality to the network load is evaluated.

Regarding the traffic engineering approaches, the objective function used for paths optimization must be formulated in an adequate, simplified way. The quality of routing decisions depends on:

- the quality of traffic prediction and traffic matrix estimation,
- the quality and granularity of the device energy models, i.e. the number of factors taken into account and the accuracy of model parameters,
- the quality of the solution – the better the solution, the more computing power must be involved (and more carbon dioxide released).

It is reasonable to assess the quality of solutions obtained by green routing algorithms using the same quality index that they use. Such approach gives us the information about the quality of the algorithms themselves, regardless of the quality of the underlying energy models of the real hardware. The assessment of energy models' quality and traffic prediction should be done separately from the assessment of algorithms, by comparing the computed optimality indexes with real-life measurement data. It should be emphasized that modelling is something individual for every single type of networking device and component. Similarly, traffic prediction covers characteristics of individual networks, and may vary with time. On the other hand, efficiency and effectiveness of optimization algorithms is something invariant – that is why it should be evaluated in separation from real-life examples.

6.2 The ECONET target energy savings

This section introduces the analytical framework for representing the impact of energy adaptive technologies and solutions for network devices under investigation by the ECONET consortium. This modelling framework is an extended version of the one proposed in [119], and aims at representing the energy gains achievable by using both standby and dynamic power scaling primitives. In case of DPS primitives, both idle logic and adaptive rate are explicitly considered.

The entire model is specifically designed for capturing the main effects and potential benefits of the energy saving primitives in a highly parametric way, without entering into the details of how these primitives will be realized inside device platforms. We modelled each energy-aware primitive by means of an extended set of efficiency and shape parameters, which allows us going beyond the simplistic representation of the device energy profile shown in Figure 85 and used in [119]. These parameters are thought to capture the main peculiarities of the energy-aware primitives under investigation. Moreover, this extended parameter set will allow us understanding how and which green primitives have to be designed in order to maximize their impact on existing network infrastructures.

Table 42: Notation used in the impact model.

Φ	The Business-As-Usual (BAU) energy consumption of the device (i.e., without any energy-aware primitives enabled)
$\Phi_{i\,min}$	The minimum value of power consumption when the network device is sleeping in the idle state
Φ_{amin}	Minimum power consumption when the network device is active
Φ_{amax}	Maximum power consumption when the network device is active
Φ_s	Power consumption when the device is fully sleeping
ψ_{ctr}	Share of energy consumption due to control plane hardware
ψ_{data}	Share of energy consumption due to data plane hardware
ψ_{cool}	Share of energy consumption due to cooling and power supply
α	Efficiency level of the standby capability
β	Efficiency level of the power scaling capability
κ	Indirect gain on cooling system and power supply
δ	Efficiency level of green traffic engineering and routing policies
S	Number of available power states
$\mu^{(s)}$	The normalized value of maximum service capacity when the device is working in the *s*-th state. Note that
λ	The normalized value of traffic load incoming to the device
ν	shape parameter of power states' energy consumption
σ	shape parameter of idle optimization efficiency (ideal case , real case)
γ	Dependency parameter of idle optimization from power states (if , idle logic is completely unaware of power states, if there is no idle logic)
$\Phi_a^{(s)}$	Maximum power consumption when the network device is active and working in the *s*-th state
$\Phi_i^{(s)}$	Power consumption when the network device is idle and working in the state
$\widetilde{\Phi}^{(s)}$	Average energy consumption of a device using LPI and AR primitives, and working in the power state
ξ	Conservative level of exploitation of green technologies
ω	Increase degree of traffic volumes in devices remaining active after the use of the green traffic engineering
$\widehat{\Phi}(\lambda)$	Average energy consumption of an active device in the presence of the incoming load
$\widetilde{\Phi}_{real}$	Average energy consumption of an active device according a traffic volume with probability density distribution
$p_\lambda(\lambda)$	Probability density distribution of incoming traffic load
$\boldsymbol{\Gamma}$	1x4 vector containing the number of device per network layer (home/access/metro/core)
$\boldsymbol{\Omega}$	1x4 vector containing the average power consumption per device per network layer (home/access/metro/core)
$\boldsymbol{\theta_d}$	1x4 vector containing the average redundancy degree of devices per network layer (home/access/metro/core)
$\boldsymbol{\theta_l}$	1x4 vector containing the average redundancy degree of links per network layer (home/access/metro/core)

Recalling the simple representation in Figure 85, we can note that the energy profile of a network device can be roughly described by means of three main parameters:

- the maximum energy consumption Φ when the device is fully active and working at the maximum speed (in such a case $\Phi = \Phi_{a\,max}$ corresponds to the energy consumption of the device with no green technologies enabled);
- the minimum energy consumption $\Phi_{a\,min}$ when the device is active but not working (in idle mode) at the minimum speed;
- the energy consumption Φ_s when the device (or parts of it) is in sleeping mode.

As already proposed by the ECONET consortium in [119], we obviously suppose that $\Phi_{\mathrm{a\,max}} \geq \Phi_{\mathrm{a\,min}} \geq \Phi_{\mathrm{s}}$; more specifically, since these optimizations will be applied at data plane HW only:

$$\Phi_{\mathrm{amax}} = \Phi \quad (11)$$

$$\Phi_{\mathrm{amin}} = \Phi[\psi_{\mathrm{ctr}} + (1-\beta)\psi_{\mathrm{data}} + (1-\kappa)\psi_{\mathrm{cool}}] \quad (12)$$

$$\Phi_{\mathrm{s}} = \Phi[\psi_{\mathrm{ctr}} + (1-\alpha)(\psi_{\mathrm{data}} + \psi_{\mathrm{cool}})] \quad (13)$$

where the ψ_* parameters represent the device internal breakdown of energy consumption sources; α, β and κ are the efficiency degrees of green technologies under investigation (the detailed notation description is in Table 42).

The way with which the energy consumption changes between the Φ and $\Phi_{i\,min}$ parameters according to the traffic load fluctuations will obviously depend on the typology, the parameterization and the efficiency of power scaling primitives (i.e., AR and LPI). In this respect, the linear relationship between workload and energy consumption in Figure 85 should certainly be considered as an ideal case, since AR and LPI may introduce different aspects that may affect system performance [120].

For instance, the intrinsic nature of AR primitives leads the device working with a finite and discrete set of stable HW configurations (in terms of input voltage, clock frequency speed, etc.), generally referred to as “state”. Depending on the HW technology in use (ASIC, FPGA, etc.) and on the implementation of AR primitives (e.g., dynamic frequency scaling - DFS, dynamic voltage scaling – DVS, dynamic frequency and voltage scaling – DVFS), the relationship between energy consumption and performance with respect to AR states may exhibit different trends (e.g., linear, quadratic, etc.).

As far as LPI primitives are concerned, we have to take the energy overhead into account, which is needed by the HW to enter and to re-wakeup from low power modes. As shown in [37][121], the more this additional energy is non-negligible, the more it makes the profile follow a curve with greater concavity with respect to Figure 85.

Owning to all these considerations, subsections 6.2.1 and 6.2.2 introduce a general model for estimating the energy profile of a device with AR and/or LPI primitives. Subsection 6.2.3 is devoted to calculate the optimal AR and LPI configurations to meet the incoming traffic load and service requirements. Subsection 6.2.4 focuses on the modelling and the usage of the sleeping/standby primitive.

6.2.1 The Adaptive Rate model

We assume the presence of S power-states. The generic s^{th} power state, with $s = 0, \ldots, S-1$, is thought to be represented by a pair composed by the maximum forwarding capacity $\mu^{(s)}$ and the maximum energy consumption $\Phi_a^{(s)}$.

The list of power states are thought to be ordered in s: s=0 corresponds to the state with the lowest energy consumption and $\mu^{(s)}$, and s=S-1 to the most energy-hungry state providing the highest network performance. In more detail, the S-1 state corresponds to the case of a device without any AR primitives enabled, and therefore $\mu^{(S-1)} = 1$ and $\Phi_a^{(S-1)} = \Phi$. For the sake of simplicity and without loss of generality, we impose that power states are equally spaced in terms of network performance:

$$\mu^{(s)} = S^{-1}(s+1) \quad (14)$$

and then $\mu^{(0)} = S^{-1}$, $\mu^{(1)} = 2\,S^{-1}$, …, $\mu^{(S-1)} = 1$.

Starting from the previous considerations, in order to estimate the maximum energy consumption in the s-th power state we apply the following relationship:

$$\Phi_a^{(s)} = \left(\Phi_{a\,min}^{\frac{1}{\nu}}\left(1-\mu^{(s)}\right) + \Phi_{a\,max}^{\frac{1}{\nu}}\,\mu^{(s)}\right)^{\nu} \quad (15)$$

It is worth noting that, in case of $\nu = 1$, we will have a linear relationship between service capacity and energy consumption, which is a typical condition of DFS approaches in many COTS processors [35], while for $\nu = 2$ we can model typical quadratic behaviour of DVS approaches in CMOS silicon. Higher values of ν could be suitable for representing DVFS approaches. Similar concepts and formulation were used also in [76].

6.2.2 The Low Power Idle model

As outlined in some previous works [121][37][120][122], the average energy consumption of a generic device with LPI primitives enabled depends on the incoming traffic volumes with respect to the system maximum capacity, and on the HW efficiency in entering and exiting from idle modes.

In order to represent such aspects, we modelled the average energy consumption $\widetilde{\Phi}^{(s)}$ of a device including idle logic and power scaling capabilities as follows:

$$\widetilde{\Phi}^{(s)}(\lambda) = \left(\left[\Phi_i^{(s)}\right]^{\frac{1}{\sigma}}\left(1-\frac{\lambda}{\mu^{(s)}}\right) + \left[\Phi_a^{(s)}\right]^{\frac{1}{\sigma}}\frac{\lambda}{\mu^{(s)}}\right)^{\sigma} \quad (16)$$

where $\Phi_i^{(s)}$ is the power consumption during idle periods, and σ the shape parameter of the power consumption curve. In detail, the case $\sigma = 1$ corresponds to the ideal case, where there is no overhead in entering and exiting idle states, and then the energy consumption increases in a linear way with respect to the traffic load λ.

As previously sketched, we expect $\Phi^{(s)}$ increasing in a convex way with respect to λ given the non-negligible idle-active transition costs. Therefore, we assume $\sigma < 1$ for real implementation of idle logic [40].

Moreover, $\Phi_i^{(s)}$ may also depend on the selected power-state:

$$\Phi_i^{(s)} = \gamma\left(\Phi_a^{(s)} - \Phi_{i\,min}\right) + \Phi_{i\,min} \quad (17)$$

where γ is the dependency degree of idle logic on the power states: if $\gamma = 0$ the energy consumption in idle periods does not depend on the s-th state. On the contrary if $\gamma = 1$, we have no idle logic optimizations in the considered device.

In more detail, and in order to provide a conservative estimation of energy saving, we decided to impose $\Phi_{i\,min} = \Phi_{a\,min}$.

6.2.3 Using the optimal energy profile with AR and/or LPI primitives

A device providing more than one power states should be controlled by optimization frameworks, working on the single device or among multiple network nodes. The specific objective of such framework is to select the best power states, providing the minimum energy consumption while meeting the incoming traffic load and QoS requirements.

From a very simplistic point of view, the output of such optimization framework can be synthesised as in the following:

$$\widehat{\Phi}(\lambda) = \min_{s=0\ldots S-1}\left\{\widetilde{\Phi}^{(s)}(\lambda)\,\middle|\,\lambda(1+\xi) \leq \mu^{(s)}\right\} \quad (18)$$

where $\xi > 0$ represents how much the optimization framework is "conservative" in allocating resources in terms of choosing the most suitable power state for meeting the incoming traffic load with the desired QoS level.

In order to make the understanding of the proposed model as intuitive as possible, Figure 116 shows an example of the energy profile according to each AR state (obtained from Eq. 6) and the final optimal one (obtained from Eq. 18) of a generic device. The number of AR profile curves is obviously equal to the number of available states (S).

The optimal profile simply consists of a piece-wise curve, composed by the most power saving parts of AR curves guaranteeing that the sum between the average incoming traffic load and a guaranteed threshold (ξ) is lower than the maximum service rate in AR states ($\mu^{(s)}$). As outlined in Figure 116, by increasing ξ, the optimal energy profile would "jump" from the AR state with lower performance to the higher one for lower values of λ. The ν parameter affects how the maximum values of AR curves are spaced among themselves. In the example of Figure 116, ν is fixed to 1 and, then, the curve joining all the maximum consumption levels of AR state profiles (i.e., $\Phi_a^{(s)}$) corresponds to a straight line; moreover, $\Phi_a^{(s)}$ values result to be equally spaced. In case that $\nu > 1$, this curve becomes more convex, and the $\Phi_a^{(s)}$ values tend to become more spaced according to the increase in s. In an analogous way, γ impacts on the energy consumption of AR state profiles when $\lambda = 0$. If γ is fixed to 0, then all the AR state profiles would start from the same level of energy consumption.

Finally, by increasing σ above 1, the curve of AR state profiles would become more concave. In case that $\sigma = 1$, the profiles become straight lines between $\Phi_a^{(s)}$ and $\Phi_i^{(s)}$.

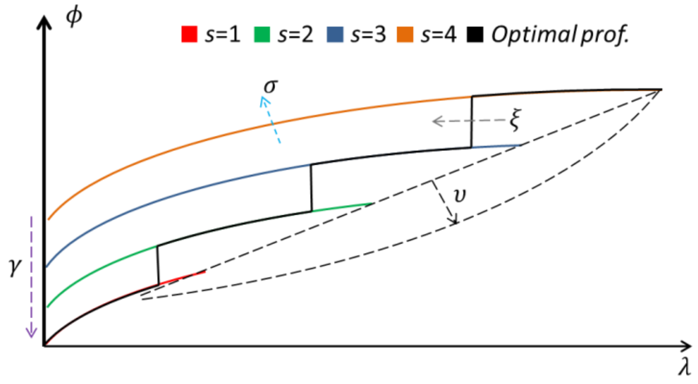

Figure 116: Example of energy profiles in various AR states in the presence of LPI primitives, and resulting optimal device energy profile.

6.2.4 Standby model

As far as the standby/sleeping primitives are concerned, the energy consumption of the entire device in standby state is obtained by means of Eq. 13.

However, we consider that they may be applied to entire devices or, depending on the prototype architecture, to selected parts of them (e.g., to a single line card, or to a link interface, etc.). However, the model here proposed does not aim at explicitly representing the energy-profile of each sub-component, but only an aggregated representation is provided.

Regarding the standby applicability, we consider to put into standby status the unused parts of the network in terms of nodes and links, which do not cause a potential “connectivity loss.” For instance, as under investigation inside the ECONET project, we consider putting into standby mode all the redundant hardware (in terms of entire backup nodes, or backup interfaces/line cards) deployed in the network for resilience purposes. In addition to redundant HW, we consider also the effect of energy-aware traffic engineering and routing policies, which may allow aggregating traffic in a subset of nodes and links during low traffic volumes periods, and consequently putting in standby status further HW across the network.

Thus, as shown also later in section 6.3, our approach leads to apply standby primitives especially on metro, transport and core network segments, where redundant HW elements and more alternative paths are generally present. At the (home) customer access level, we suppose to use only LPI primitives.

6.2.5 Considering the entire network infrastructure

In a similar way with respect to the work in [119], we consider a network infrastructure composed by 4 network segments: home, access, metro and transport, and core. The representative devices for such network levels are home gateways, Digital Subscription Line Access Multiplexers (DSLAMs), high-end optical switches, and multi-chassis IP routers, respectively. Thus, in order to represent such network structure in a compact way, we will use a vector notation. The vectors and matrixes will be represented by using the bold style.

We start by introducing the $\mathbf{\Gamma}$ and $\mathbf{\Omega}$ vectors, having 4 elements representing the device density and related BAU power consumption (Φ) for the home, access, metro/transport and core networks, respectively. Thus, BAU energy requirement of the entire network infrastructure can be easily expressed as:

$$\mathrm{E_{BAU}} = \mathbf{\Gamma} \cdot \mathbf{\Omega} \quad (19)$$

In order to estimate the energy gain of green technologies, we need to estimate which share of HW (in terms of entire or part of them) can enter standby mode, and consequently how much HW is active and working with LPI and AR.

To this purpose, we can simply obtain the share of redundant HW (to be put in standby mode) as follows:

$$\boldsymbol{\theta}_{\mathbf{red}} = \frac{\boldsymbol{\theta}_{\mathbf{d}}}{2} + \left(1 - \frac{\boldsymbol{\theta}_{\mathbf{d}}}{\mathbf{2}}\right) \mathbf{I_4}\, \frac{\boldsymbol{\theta}_{\mathbf{l}}}{2} \quad (20)$$

where $\mathbf{I_4}$ is the identity matrix of rank 4, while $\boldsymbol{\theta}_{\mathbf{d}}$ and $\boldsymbol{\theta}_{\mathbf{l}}$ are the average redundancy degree of devices and link, respectively.

As previously sketched, the adoption of green traffic engineering and routing policies may help in aggregating the transport traffic flows in a subset of nodes and links, especially during night hours with low traffic volumes. So, additional HW that can enter standby modes can be estimated as follows:

$$\boldsymbol{\theta}_{\boldsymbol{\alpha}} = \boldsymbol{\theta}_{\mathbf{red}} + (1 - \boldsymbol{\theta}_{\mathbf{red}})\, \mathbf{I_4}\, \boldsymbol{\delta} \quad (21)$$

On the other hand, aggregating the traffic in few nodes and links causes an increase in traffic load on the ones remaining active. The increase factor of traffic volumes in devices remaining active can be expressed as the ratio of active network resources without such policies, and the number of active ones with the policies:

$$\boldsymbol{\omega}^{(j)} = \frac{1-\boldsymbol{\theta}_{\boldsymbol{red}}{}^{(j)}}{1-\boldsymbol{\theta}_{\boldsymbol{\alpha}}{}^{(j)}} \quad \text{with } j = 0, \dots, 3 \quad (22)$$

where j represents the j-th element of $\boldsymbol{\omega}$, $\boldsymbol{\theta}_{\boldsymbol{red}}$ and $\boldsymbol{\theta}_{\boldsymbol{\alpha}}$ vectors.

Extending Eq. 19, the energy consumption of the entire network infrastructure working with green technologies enabled can be easily expressed as the energy consumption of sleeping HW elements and active ones:

$$\mathrm{E}_{GT} = \left[\boldsymbol{\theta}_{\boldsymbol{\alpha}}\, \mathbf{I_4} \boldsymbol{\Phi}_{\boldsymbol{s}} + (\mathbf{1} - \boldsymbol{\theta}_{\boldsymbol{\alpha}})\, \mathbf{I_4}\, \widetilde{\boldsymbol{\Phi}}_{\boldsymbol{real}}\right] \cdot \mathbf{\Gamma} \quad (23)$$

where $\widetilde{\boldsymbol{\Phi}}_{\mathbf{real}}$ is the average power consumption (per network segment) of active devices working with AR and LPI in the presence of fluctuating incoming traffic loads.

In more detail, in order to capture the effect of traffic daily profiles shown in Section 4 and recalling Eq. 18, we express $\widetilde{\Phi}_{\mathrm{PS}}$ as the expectation of $\widehat{\Phi}(\lambda\omega)$ with respect to the incoming traffic volume λ and the increase of traffic ω due to the additional resources put in standby mode by network control policies:

$$\widetilde{\boldsymbol{\Phi}}_{\boldsymbol{PS}} = \int_0^1 \widehat{\boldsymbol{\Phi}}(\lambda \boldsymbol{\omega})\, p_\lambda(\lambda) d\lambda \quad (24)$$

Finally, the energy saving with respect to the BAU case can be easily found as follows:

$$\Delta\mathrm{E_{GT}} = \mathrm{E_{BAU}} - \mathrm{E_{GT}} \quad (25)$$

6.3 Estimating the energy saving gains

This section aims at evaluating the potential impact estimated by the proposed model within the energy-efficiency degrees suggested by the ECONET consortium in the presence of the realistic scenarios provided by Telecom Italia and GRNET for the access, home, metro/transport, and core networks. The main aim of this analysis is to understand the impact of each energy adaptation primitive on the various network segments, and consequently to try collecting precious information on how these technologies have to be designed and which level of efficiency they should provide for maximizing energy consumption reduction in real-world ISP and Telco network infrastructures.

As far as the network scenarios are concerned, we fixed the Γ and Ω for GRNET and Telecom Italia by starting from Table 15 and Table 17:

$$\Gamma_{\text{TELIT}} = \begin{vmatrix} 17500000 \\ 27344 \\ 1750 \\ 175 \end{vmatrix} \quad \Omega_{\text{TELIT}} = \begin{vmatrix} 10 \\ 1280 \\ 6000 \\ 10000 \end{vmatrix} \quad \Gamma_{\text{GRNET}} = \begin{vmatrix} 0 \\ 42 \\ 0 \\ 11 \end{vmatrix} \quad \Omega_{\text{GRNET}} = \begin{vmatrix} 0 \\ 297 \\ 0 \\ 3677 \end{vmatrix} \quad (26)$$

It is worth noting that the network access segment of Telecom Italia (the 2nd elements of the vectors above) has a quite different nature in terms of device typologies and technologies from the GRNET one, since, for the former, it corresponds to the residential access, whereas, for the latter, it corresponds to peering to local academic networks, research institutes or data-centers. In terms of device typology (mainly L2 switches), the access segment of GRNET roughly corresponds to the metro one of Telecom Italia.

Starting from the values in Table 16and Table 18, we can define the redundancy degrees per network level in both the considered scenarios:

$$\theta_{d|\text{TELIT}} = \begin{vmatrix} 0.00 \\ 0.00 \\ 0.13 \\ 1.00 \end{vmatrix} \quad \theta_{d|\text{TELIT}} = \begin{vmatrix} 0.00 \\ 0.88 \\ 0.00 \\ 0.72 \end{vmatrix} \quad \theta_{l|\text{TELIT}} = \begin{vmatrix} 0.00 \\ 0.00 \\ 1.00 \\ 0.50 \end{vmatrix} \quad \theta_{l|\text{GRNET}} = \begin{vmatrix} 0.00 \\ 0.13 \\ 0.00 \\ 0.71 \end{vmatrix} \quad (27)$$

Regarding the traffic profiles, we used the data obtained by real measurements in the GRNET and TELIT networks. In more detail, the traffic load measurements were used to obtain an estimation of the load probability distribution $p_\lambda(\lambda)$, while, in order to make a conservative estimation, the maximum link occupancy was fixed at 50% of the available capacity.

As far as the devices in the various network segments are concerned, we use the same energy consumption breakdown as in [119], and coming from device manufacturers in the ECONET consortium:

$$\psi_{\text{ctr}} = \begin{vmatrix} 0.03 \\ 0.03 \\ 0.13 \\ 0.11 \end{vmatrix} \quad \psi_{\text{data}} = \begin{vmatrix} 0.95 \\ 0.80 \\ 0.73 \\ 0.54 \end{vmatrix} \quad \psi_{\text{cool}} = \begin{vmatrix} 0.02 \\ 0.17 \\ 0.14 \\ 0.35 \end{vmatrix} \quad (28)$$

In a similar way, also regarding the energy-efficiency parameters, we used the values indicated by the ECONET consortium:

$$\alpha = \begin{vmatrix}0.8\\0.8\\0.8\\0.8\end{vmatrix} \quad \beta = \begin{vmatrix}0.7\\0.7\\0.7\\0.7\end{vmatrix} \quad \kappa = \begin{vmatrix}.15\\.15\\.15\\.15\end{vmatrix} \quad \gamma = \begin{vmatrix}0\\0\\0\\0\end{vmatrix} \quad s = \begin{vmatrix}5\\5\\5\\5\end{vmatrix} \quad \sigma = \begin{vmatrix}0.5\\0.5\\0.5\\0.5\end{vmatrix} \quad \delta = \begin{vmatrix}0\\0\\0.2\\0.2\end{vmatrix} \quad \upsilon = \begin{vmatrix}2\\2\\2\\2\end{vmatrix} \quad \xi = \begin{vmatrix}0.1\\0.1\\0.1\\0.1\end{vmatrix} \quad (29)$$

In the rest of this section it is estimated the energy-saving that the green technologies can provide in the Telecom Italia and GRNET network infrastructures if implemented according to the efficiency parameters defined in the ECONET project. Then, this is compared with the “Business As Usual” scenarios in order to understand the benefits per each network segment, and for each green primitive under investigation.

6.3.1 Overall energy consumption

Considering the TELIT and GRNET reference scenarios described in Section 4.2, we compute the overall energy consumption as element of comparison for home, access, metro/transport, and core networks.

The yearly energy consumption in the BAU scenario is estimated with Eq. 19 using the energy consumption and the number of devices of the two networks shown in Eq. 26. Figure 117 shows the BAU yearly energy consumption for home, access, metro/transport, and core networks. The most of the energy consumption for the TELIT reference scenario is in the home network. Instead, the GRNET reference scenario does not include the home and metro/transport networks, and the most energy consumption is for the core one.

Using the Eq. 23, we compute the energy consumption estimation through the proposed energy profile model. Figure 118 shows the energy saving gains between the energy consumption estimated by the model including the Dynamic Power Scaling (DPS) and Stand-by primitives and the BAU scenario. It is easy to see that there are not substantial differences between the working day and holiday profiles. In the TELIT scenario, the most energy saving is in the metro/transport and core networks; while the minimum one is in the Access one. Instead, for GRNET, also the access network which has an high value of energy saving gain. As already described, GRNET is a medium size ISP with the access network that can be compared to a metro/transport network of a large-scale telco like TELIT (obviously with different sizes).

In order to evaluate the behaviour of the model we estimate the energy saving gain enabling only DPS primitives or the Stand-by ones. Figure 119 and Figure 120 report these gains for TELIT and GRNET traffic traces, respectively.

It is interesting to outline that for the Telecom Italia reference scenario, the DPS primitives contribute to the energy saving gain in every part of the network, while the stand-by ones yield a significant energy saving gain only for metro/transport and core networks. This effect is simply due to the fact that we suppose standby capabilities to be applied only where “alternative paths” are present, and then only to those segments with partially/fully meshed topologies. Obviously, residential access generally has no such features.

As a result, standby primitives seem to assure major gains in network segments with high meshing degrees, and DPS ones in the other cases.

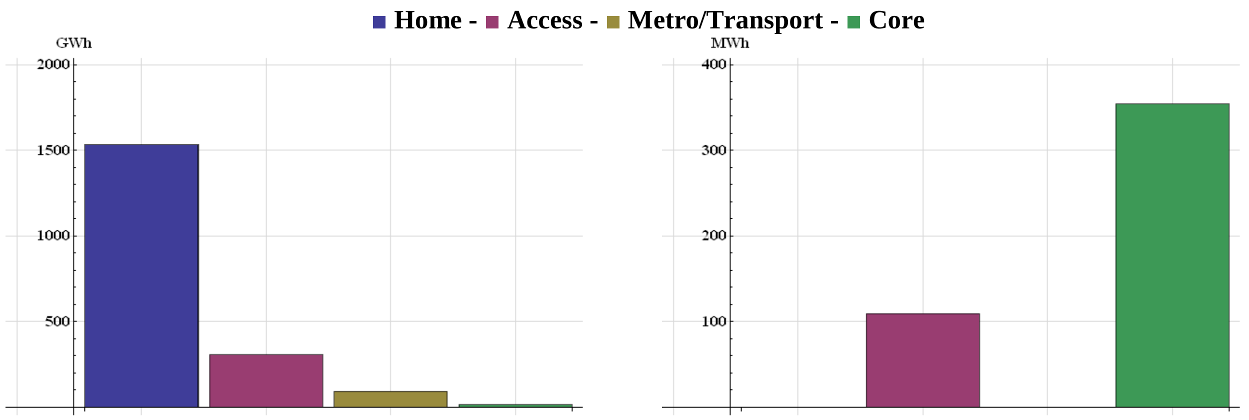

(*a*) TELIT (*b*) GRNET

Figure 117: Yearly Energy consumption estimation for TELIT (*a*) and GRNET (*b*) Home, Access, Metro/Transport, and Core networks in the BAU scenario.

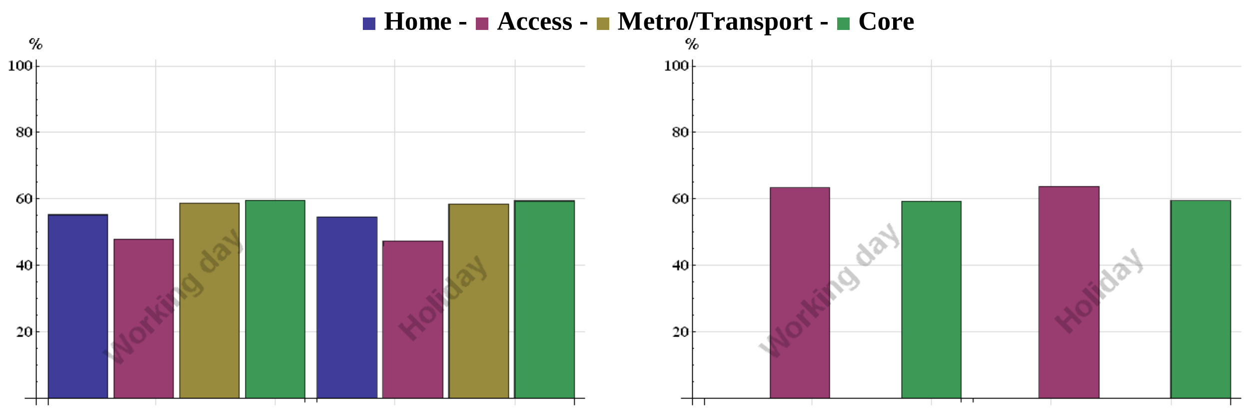

(*a*) TELIT (*b*) GRNET

Figure 118: Energy saving estimation for TELIT (*a*) and GRNET (*b*) Home, Access, Metro/Transport, and Core networks during working day and holiday profiles between the proposed energy model and the BAU scenarios.

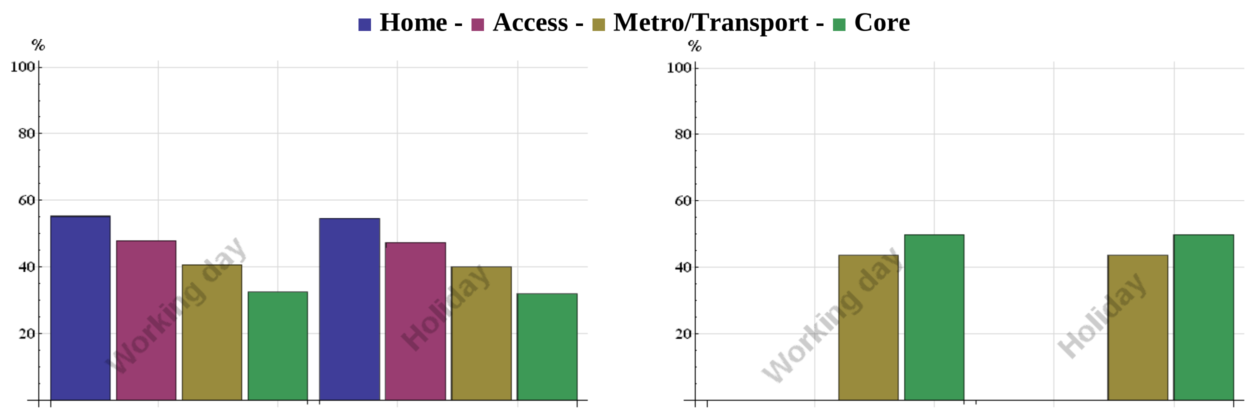

(*a*) TELIT (DPS primitives) (*b*) TELIT (Stand-by primitives)

Figure 119: Energy saving estimation for TELIT Home, Access, Metro/Transport, and Core networks during working day and holiday profiles between the proposed energy model with only DPS (*a*) or Stand-by (*b*) primitives and the BAU scenarios.

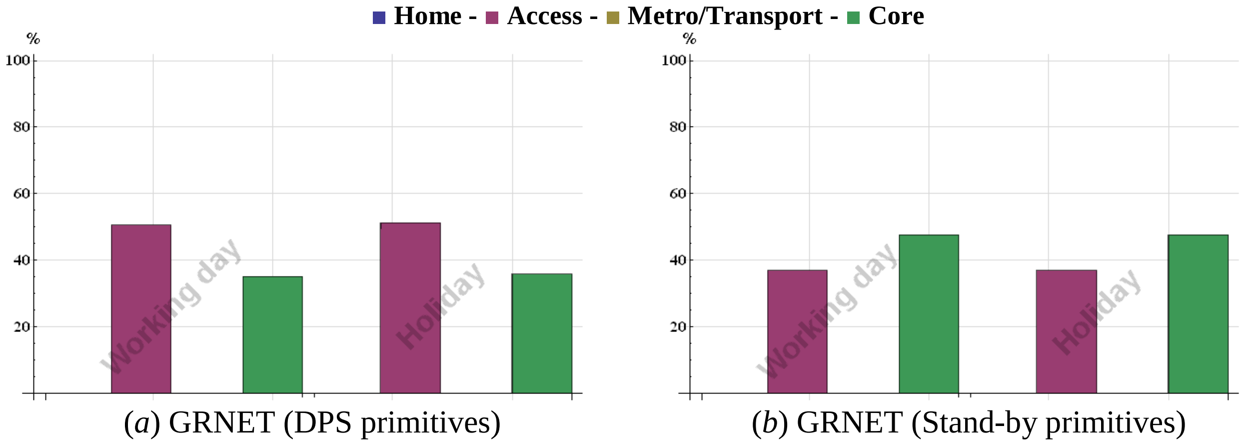

Figure 120: Energy saving estimation for GRNET Home, Access, Metro/Transport, and Core networks during working day and holiday profiles between the proposed energy model with only DPS (*a*) or Stand-by (*b*) primitives and the BAU scenarios.

6.4 Summary

In section 6, general performance indexes have been defined for the evaluation of the mechanisms that will be developed within the project. The section initially provided a brief state of the art of the existing standards, in order to point out their strongest and weakest points. Based on the description of the standards, the benchmarking methodology that will be followed in the project were detailed. The methodology that will be followed within ECONET consists in determining two indexes: one represents the amount of energy that can be saved through power management capabilities, while the other one quantifies the performance degradation that these capabilities can cause.

Specific performance indexes and evaluation criteria were defined for both the device data and the control plane, aiming to denote the trade-off between power consumption and network performance. In the data plane, the metrics include the power saving obtained thanks to power management in comparison to a scenario with no such capabilities, and the performance degradation as a ratio between the values of packet latency in the case of an ideal network device and the ones measured with the real system under test. In the control plane, evaluation criteria include the level of energy proportionality to the network load the equipment may reach. This comparison can be realized through the estimation of the average energy in Watts required for serving 1 bit (average energy-per-bit). Furthermore, the Perceptual QoS (PQoS) – or Quality of Experience (QoE) – is defined as a metric of the quality experienced by the end user. PQoS has to be evaluated against classic QoS and achieved energy savings.

Finally, an analytical framework was presented for representing the impact of energy adaptive technologies and solutions for network devices under investigation by the ECONET consortium. Based on this framework, an evaluation was conducted for the potential impact estimated by the proposed model within the energy-efficiency degrees suggested by the ECONET consortium in the presence of the realistic scenarios provided by Telecom Italia and GRNET for the access, home, metro/transport, and core networks. Based on this analysis, it was shown that energy savings up to more than 60% can be realized in the different network parts.

7 Conclusions

Since ECONET is devoted at re-thinking and re-designing wired network equipment and infrastructures towards more energy-sustainable and eco-friendly technologies and perspectives, in this deliverable, the main energy aware technologies that will be examined in the project have been presented. These technologies were initially mapped to the energy aware design space and analysed in detail per network part for the data and control plane. Specific requirements were extracted that will guide the development of green technologies for the data plane of network devices (WP3), the definition of the green abstraction layer for exposing them to higher level protocols (WP4), as well as the design of the control plane strategies running higher up at the protocol level (WP5).

In the first part of the deliverable, the necessity has been claimed for the design and introduction of novel and energy-aware techniques at the device and network level and for the reduction in the energy consumption of network equipment with specific targets in the short and long run. This necessity is based on the fact that the ICT sector accounts for more than 2% of global carbon emissions with an increasing trend and predictions for the following years. This necessity is also amplified by taking into account the increasing trends in the traffic volumes exchanged within the Internet and the energy consumption disproportionality that are present in today’s networking devices.

Energy consumption of current devices under variable conditions have been measured and presented and specific requirements that have to be fulfilled from improving the devices’ energy efficiency have been described for the home, access, metro/transport and core networking parts. Based on these requirements, a target energy saving has been defined per device type with energy gains ranging from 20% to 70% based on the applied technology in each device part. In addition to the definition of requirements per device type, three reference scenarios have been described in detail aiming at the examination of the network topologies and traffic profiles in telecom operators and Internet service provider networks and at the extraction of conclusions regarding the energy gain that can be attained through the application of network wide energy aware techniques. It has been shown that such a margin exists due to the overprovisioning of current network links, the existence of multiple alternative paths in the core networking part and the variations spotted in the various links’ traffic profiles (especially the day and night variations). For each scenario, the main technologies that will be examined within ECONET have been presented.

In the second part of the deliverable, an in-depth analysis of the energy aware technologies that can be applied in the data and control plane has been presented. In the network device data plane we described re-engineering approaches aiming at introducing and designing more energy-efficient elements for network device architectures, dynamic adaptation mechanisms for modulating capacities of network modules to meet actual traffic loads and requirements, and sleeping/standby approaches for smartly and selectively driving unused network/device portions to low standby modes. In the control plane we analysed energy aware traffic engineering strategies targeting to achieve effectively serving traffic while reducing the network’s overall energy footprint, virtualization techniques for achieving more efficient utilization of resources, and resource allocation and algorithmic efficiency mechanisms directly affecting energy consumption and providing the potential for further energy saving.

Based on the defined requirements and the description of the energy-aware technologies, energy consumption profiles have been provided for the prototype devices that will be examined within ECONET. It has been shown that by applying energy-aware techniques in various types of network devices (routers, switches, home gateways, FPGA based systems), reduction in energy consumption up to 60% or even 80% may be accomplished. Monitoring and management issues were also examined with regard to energy consumption metrics and the basic energy-related parameters that have to be monitored were defined.

Finally, the evaluation methodology that will be followed within the project has been detailed. We described specific methods and performance metrics for the evaluation of the developed mechanisms in the data and control plane. The impact of the application of specific energy aware techniques on the potential energy saving has been examined in overall networking infrastructures under the realistic scenarios provided by TELIT, GRNET and NASK. It has been shown that through the design and application of proper energy aware techniques, it is possible to achieve more than 50% of energy saving in short-term.

References

[1] C. Bianco, F. Cucchietti, G. Griffa, "Energy consumption trends in the next generation access network — a telco perspective," Proc. 29th Internat. Telecommunications Energy Conf. (INTELEC 2007), Rome, Italy, Sept. 2007, pp.737-742

[2] R. Bolla, R. Bruschi, F. Davoli, F. Cucchietti, "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures," IEEE Communications Surveys and Tutorials (COMST).

[3] Telecom Italia Website, "The Environment," URL: http://www.telecomitalia.it/sostenibilita2006/ English/B05.html

[4] British Telecom Group, "Sustainability Report 2009", URL: http://www.btplc.com/Societyandenvironment/Ourapproach/Sustainabilityreport/index.aspx.

[5] British Telecom Group, "Sustainability Report 2011", URL: http://www.btplc.com/Responsiblebusiness/Ourstory/Sustainabilityreport/Our2011sustainabilityreview/index.aspx.

[6] BT Press, "BT announces major wind power plans," 2007, URL: http://www.btplc.com/News/Articles/Showarticle.cfm?ArticleID=dd615e9c-71ad-4daa-951a-55651baae5bb.

[7] ITWales, "Green evangelist to call for big changes in computer use to aid environment at ITWales Conf.," Nov. 2007, http://www.itwales.com/997539.htm.

[8] Deutsche Telekom Group, "Ecological Indicators", URL: http://cr-report.telekom.de/site08/en/daten-fakten/kennzahlen/oekologische-kennzahlen-/index.php.

[9] S. N. Roy, "Energy Logic: A Road Map to Reducing Energy Consumption in Telecommunications Networks," Proc. 30th Internat. Telecommunications Energy Conference (INTELEC 2008), San Diego, CA, USA, Sept. 2008.

[10] Annual Corporate Responsibility Report Telefonica, S.A., 2009,

[11] H. Ikebe, N. Yamashita, R. Nishii, "Green Energy for Telecommunications," Proc. 29th Internat. Telecommunications Energy Conf. (INTELEC 2007), Rome, Italy, Sept. 2007, pp. 750-755.

[12] European Commission DG INFSO, "Impacts of Information and Communication Technologies on Energy Efficiency," final report, Sept. 2008, ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/sustainable-growth/ict4ee-final-report_en.pdf.

[13] U.S. Energy Information Administration (EIA), Official Energy Statistics from the U.S. Governement, website http://www.eia.doe.gov.

[14] Global e-Sustainibility Initiative (GeSI), "SMART 2020: Enabling the low carbon economy in the information age", http://www. theclimategroup.org/assets/resources/publications/Smart2020Report.pdf.

[15] H. S. Matthews, W. S. Loh, H. M. Chong, “Electricity Use of Wired and Wireless Telecommunications Networks in the United States” Proc. IEEE Internat. Symp. on Electronics and the Environment, Boston, MA, USA, May 2003, pp. 131-136, 2003

[16] M. Etoh, T. Ohya, Y. Nakayama, “Energy Consumption Issues on Mobile Network Systems,” Proc. Internat. Symp. on Applications and the Internet (SAINT 2008), Turku, Finland, 2008, pp. 365-368.

[17] R. S. Tucker, R. Parthiban, J. Baliga, K. Hinton, R. W. A. Ayre, W. V. Sorin, “Evolution of WDM Optical IP Networks: A Cost and Energy Perspective”, IEEE Journal of Lightwave Technology, vol. 27, no. 3, pp. 243-252, Feb 2009.

[18] R. S. Tucker, “Optical packet-switched WDM networks-a cost and energy perspective,” Proc. Opt. Fiber Commun. Conf. Expo. (OFC 2008), San Diego, CA, USA, March 2008.

[19] J. Baliga, R. Ayre, K. Hinton, and R. S. Tucker, “Photonic switching and the energy bottleneck,” Proc. Internat. Conf. Photonics in Switching 2007, San Francisco, CA, 2007.

[20] J. Baliga, K. Hinton, and R. S. Tucker, “Energy consumption of the Internet,” Proc. COIN-ACOFT, Melbourne, Australia, 2007.

[21] J. Baliga, R. Ayre, W. Sorin, K. Hinton, and R. S. Tucker, “Energy consumption in access networks,” Proc. Optical Fiber Commun. Conf. Expo, San Diego, CA, March 2008.

[22] Cisco CRS-1 Carrier Routing System, http://www.cisco.com/en/US/products/ps5763/index.html.

[23] Juniper T1600 Core Router, http://www.juniper.net/us/en/products-services/routing/t-tx-series/t1600/.

[24] Huawei, Qudway NetEngine 5000E Core Router, http://www.huawei.com/products/datacomm/catalog.do?id=25.

[25] Brocade NetIron XMR 160000 and 320000 Internet backbone routers, http://www.brocade.com/sites/dotcom/products-solutions/products/ ethernet-switches-routers/service-provider/product-details/netiron-xmr-series/index.page.

[26] The Energy Consumption Rating (ECR) Initiative, ECR Specification v3.0.1, http://www.ecrinitiative.org/pdfs/ECR_3_0_1.pdf

[27] Neilson, D.T., "Photonics for switching and routing," IEEE Journal of Selected Topics in Quantum Electronics (JSTQE), vol.12, no.4, pp.669-678, July-Aug. 2006.

[28] M. Bohr, “A 30 Year Retrospective on Dennard’s MOSFET Scaling Paper,” IEEE SSCS Newsletter, vol. 12, no. 1, pp. 11-13, Winter 2007.

[29] L. Ceuppens, A. Sardella, D. Kharitonov, “Power Saving Strategies and Technologies in Network Equipment Opportunities and Challenges, Risk and Rewards,” Proc. Internat. Symp. on Applications and the Internet (SAINT 2008), Turku, Finland, July 2008, pp. 381-384.

[30] Juniper “T Series Core Routers Architecture Overview - High-End Architecture for Packet Forwarding and Switching,” White Paper, 2010, http://www.juniper.net/us/en/local/pdf/whitepapers/2000302-en.pdf

[31] L. G. Roberts, “A Radical New Router”, IEEE Spectrum, vol. 46, no. 7, pp. 34-39, July 2009.

[32] M. Baldi, Y. Ofek, “Time for a “Greener” Internet,” Proc. Green Communications Workshop in conjunction with IEEE ICC'09 (GreenComm09), Dresden, Germany, June 2009.

[33] D. Harris, “Reducing FPGA Power Consumption”, Electronic Design, Penton Media Inc., Feb. 2008.

[34] J. Noguera, I.O. Kennedy, “Power Reduction in Network Equipment Through Adaptive Partial Reconfiguration,” Proc. 2007 Internat. Conf. on Field Programmable Logic and Applications (FPL 2007), Aug. 2007, pp. 240-245.

[35] R. Bolla, R. Bruschi, F. Davoli, A. Ranieri, “Performance Constrained Power Consumption Optimization in Distributed Network Equipment,” Proc. Green Communications Workshop in conjunction with IEEE ICC'09 (GreenComm09), Dresden, Germany, June 2009.

[36] A. Wierman, L. L. H. Andrew, A. Tang, “Power-Aware Speed Scaling in Processor Sharing Systems,” Proc. of the 28th IEEE Conference on Computer Communications (INFOCOM 2009), Rio de Janeiro, Brazil, Apr. 2009.

[37] S. Nedevschi, L. Popa, G. Iannaccone, D. Wetherall, S. Ratnasamy, “Reducing network energy consumption via sleeping and rate-adaptation,” Proc. 5th USENIX Symp. on Networked Systems Design and Implementation (NSDI '08), San Francisco, CA, 2008, pp. 323-336.

[38] M. B. Srivastava, A.P. Chandrakasan, R.W. Brodersen, “Predictive System Shutdown and Other Architectural Techniques for Energy Efficient Programmable Computation,” IEEE Trans. Very Large Scale Integration (VLSI) Systems, vol. 4, no. 1, pp. 42-55, Mar. 1996.

[39] L. Benini, A. Bogliolo, G. Paleologo, and G. De Micheli, “Policy Optimization for Dynamic Power Management,” IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems, vol. 18, no. 6, pp. 813-833, June 1999.

[40] E. Y. Chung, L. Benini, A. Bogliolo, Y.-H. Lu, and G. De Micheli, “Dynamic Power Management for Nonstationary Service Requests,” IEEE Trans. Computers, vol. 51, no. 11, pp. 1345-1361, Nov. 2002.

[41] A. Miyoshi, C. Lefurgy, E. Van Hensbergen, R. Rajamony, R. Rajkumar, “Critical power slope: understanding the runtime effects of frequency scaling,” Proc. 16th Internat. Conf. on Supercomputing (ICS’02), New York, NY, USA, June 2002.

[42] Z. Ren, B. Krogh, and R. Marculescu, “Hierarchical Adaptive Dynamic Power Management,” IEEE Trans. Computers, vol. 54, no. 4, pp. 409-420, Apr. 2005.

[43] R. Golding, P. Bosch, and J. Wilkes, “Idleness Is Not Sloth,” Technical Report HPL-96-140, Hewlett-Packard Laboratories, Oct. 1996.

[44] S. Gochman et al., “The Intel Pentium M Processor: Microarchitecture and Performance,” Intel Technology J., vol. 7, no. 2, pp. 21- 36, May 2003.

[45] S. Herbert, D. Marculescu, “Analysis of dynamic voltage/frequency scaling in chip-multiprocessors,” Proc. 2007 Internat. Symp. on Low Power Electronics and Design (ISLPED '07), Portland, OR, USA, Aug. 2007, pp. 38-43.

[46] R. Chabukswar, “Maximizing Performance and Energy-Efficiency on Intel@ CoreTM Microarchitecture using Multi-Threading”, Intel White paper.

[47] K. Christensen, C. Gunaratne, B. Nordman, and A. George, “The Next Frontier for Communications Networks: Power Management,” Computer Comm., vol. 27, no. 18, pp. 1758-1770, Dec. 2004.

[48] K. Christensen, C. Gunaratne, B. Nordman, “Managing Energy Consumption Costs in Desktop PCs and LAN Switches with Proxying, Split TCP Connections, and Scaling of Link Speed,” Internat. Journal of Network Management, vol. 15, no. 5, pp. 297-310, Sept./Oct. 2005.

[49] M. Allman, K. Christensen, B. Nordman, V. Paxson, “Enabling an Energy-Efficient Future Internet Through Selectively Connected End Systems,” Proc. ACM SIGCOMM HotNets Workshop (HotNets 07), Atlanta, GA, Nov. 2007.

[50] M. Jimeno, K. Christensen, B. Nordman, “A Network Connection Proxy to Enable Hosts to Sleep and Save Energy,” Proc. IEEE Internat. Performance Computing and Communications Conf., pp. 101-110, Dec. 2008.

[51] [Online]: Ganglia Monitoring System - http://ganglia.sourceforge.net/

[52] [Online]: AMEE - http://my.amee.com/developers

[53] The State of the Internet, Volume 3, Number 2, 2nd Quarter, 2010.

[54] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, F. Jahanian, “Internet inter-domain traffic”, SIGCOMM Comput. Commun. Rev. 40, 4, 75-86, August 2010.

[55] OECD Broadband Portal, http://www.oecd.org/sti/ict/broadband .

[56] G. Maier, A. Feldmann, V. Paxson, M. Allman, “On dominant characteristics of residential broadband internet traffic”, In Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference (IMC '09), 2009.

[57] Cisco Visual Networking Index: Forecast and Methodology, 2010-2015.

[58] K. Sato, "Optical technologies that enable Green networks," 12th International Conference on Transparent Optical Networks (ICTON), pp.1-4, June 27 2010-July 1 2010

[59] R. S. Tucker, "Green Optical Communications---Part II: IEEE Journal of Energy Limitations in Networks," Selected Topics in Quantum Electronics, no.99, pp.1-14, 2010.

[60] R. Bolla, R. Bruschi, "An effective forwarding architecture for SMP Linux routers," Telecommunication Networking Workshop on QoS in Multiservice IP Networks, 2008. IT-NEWS 2008. 4th International , Venice, Italy, Feb. 2008, pp.210-216.

[61] R. Bolla, R. Bruschi, “PC-based Software Routers: High Performance and Application Service Support,” Proc. of the ACM Sigcomm Worshop on Programmable Routers for Extensible Services of Tomorrow (PRESTO’08), Seattle, WA, USA, Aug. 2008.

[62] Intel, "Intel is Leading the Way in Designing Energy-Efficient Platforms", Intel White paper.

[63] The Advanced Configuration & Power Interface (ACPI) standard, version 5.0, http://www.acpi.info/.

[64] M. Xia, M. Tornatore, Y. Zhang, P. Chowdhury, C. Martel, B. Mukherjee, "Greening the Optical Backbone Network: A Traffic Engineering Approach," Communications (ICC), 2010 IEEE International Conference on , vol., no., pp.1-5, 23-27 May 2010.

[65] U. Lee, I. Rimac, V. Hilt, "Greening the internet with content-centric networking", In Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking (e-Energy '10). ACM, New York, NY, USA, 179-182, 2010.

[66] P. Mell, T. Grance, Draft NIST Working Definition of Cloud Computing v14, Nat. Inst. Standards Technol., 2009. [Online]. Available: http://csrc.nist.gov/groups/SNS/cloud-computing/index.html.

[67] J. Baliga, R.W.A. Ayre, K. Hinton, R.S. Tucker, "Green Cloud Computing: Balancing Energy in Processing, Storage, and Transport," Proceedings of the IEEE , vol.99, no.1, pp.149-167, Jan. 2011

[68] N. Vasic, D. Kostic, "Energy-aware traffic engineering", 1st International Conference on Energy-Efficient Computing and Networking (e-Energy '10), ACM, New York, USA, 2010.

[69] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, S. Wright, "Power Awareness in Network Design and Routing", IEEE INFOCOM 2008.

[70] Juniper, Energy Efficiency for Network Equipment: Two Steps Beyond Green Washing, White paper, 2010, http://www.juniper.net/us/en/local/pdf/whitepapers/2000284-en.pdf

[71] N. Vasic, D. Novakovic, S. Shekhar, P. Bhurat, M. Canini, D. Kostic, "Responsive Energy-Proportional Networks", EPFL Technical Report, EPFL-REPORT-149848, 2010.

[72] M. Gupta, S. Singh, "Greening of the Internet", ACM SIGCOMM, 2003.

[73] D. Rossi, A. P. Bianzino, J.-L. Rougier, C. Chaudet, and F. Larroca, "Energy-Aware Routing: a Reality Check", GLOBECOM Green Communications Workshop '10, Miami, FL, Dec. 2010.

[74] M. Zhang, C. Yi, B. Liu, B. Zhang, "GreenTE: Power-Aware Traffic Engineering", IEEE International Conference on Network Protocols (ICNP), 2010.

[75] A.Cianfrani, M.Listanti, V.Eramo, M.Marazza, E.Vittorini, "An energy saving routing algorithm for a green OSPF protocol," Proc. IEEE Infocom'10, San Diego, CA, Mar. 2010.

[76] J.Restrepo, C.Gruber, C.Machoca, "Energy Profile Aware Routing," Proc. IEEE GreenComm'09, Dresden, Germany, June 2009.

[77] B.Nordman, K.Christensen, "Proxying: The Next Step in Reducing IT Energy Use," IEEE Computer, Vol. 43, No. 1, pp. 91-93, January 2010.

[78] Y.Wang, E.Keller, B.Biskeborn, J.van der Merwe, J.Rexford, "Virtual Routers on the Move: Live Router Migration as a Network-Management Primitive," Proc. ACM SIGCOMM, Seattle, WA, Aug. 2008.

[79] VMware, Virtualization overview, White paper, 2006, [Online]: http://www.vmware.com/pdf/virtualization.pdf

[80] VMware, Understanding Full Virtualization, Paravirtualization, and Hardware Assist, White paper, 2007, [Online]: http://www.vmware.com/files/pdf/VMware_paravirtualization.pdf

[81] T. Abels, P. Dhawan, B. Chandrasekaran, "An Overview of Xen Virtualization", Dell power solutions, August 2005, [Online]: http://www.dell.com/downloads/global/power/ps3q05-20050191-Abels.pdf

[82] L. A. Barroso, U. Hölzle, "The case of energy-proportional computing", IEEE Computer 40(12) (2007) pp. 33-37

[83] VMware, VMware TCO / ROI Methodology, White paper, 2007, [Online]: http://www.vmware.com/files/pdf/tco_roi_methdology.pdf

[84] R. Fichera, "Power And Cooling Heat Up The Data Center", Forrester Research, March 8, 2006

[85] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, A. Warfield, "Live Migration of Virtual Machines", NSDI 2005

[86] T. Das, P. Padala, V. N. Padmanabhan, "LiteGreen: Saving Energy in Networked Desktops Using Virtualization", USENIX 2010

[87] Y. Agarwal, S. Savage, R. Gupta, "SleepServer: A Software-Only Approach for Reducing the Energy Consumption of PCs within Enterprise Environments", USENIX ATC 2010

[88] DMTF, DSP2017 1.0.0 Open Virtualization Format, White Paper, 2009, http://www.dmtf.org/sites/default/files/standards/documents/DSP2017_1.0.0.pdf

[89] A. Beloglazov, R. Buyya, "Energy Efficient Resource Management in Virtualized Cloud Data Centers", CCGrid 2010

[90] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag and B. Maggs, "Cutting the Electric Bill for Internet-Scale Systems", SIGCOMM 2009

[91] J. Reich, M. Goraczko, A. Kansal, J. Padhye, "Sleepless in Seattle No Longer", USENIX 2007

[92] R. Bianchini, R. Rajamony, "Power and energy management for server systems", IEEE Computer 37(11) (2004) pp. 68-76

[93] Cisco, Router Virtualization in Service Providers, White paper, 2008, http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns562/ns573/white_paper_c11-512753.pdf

[94] Juniper, Control Plane Scaling and router Virtualization, White paper, 2010, http://www.juniper.net/us/en/local/pdf/whitepapers/2000261-en.pdf

[95] FEDERICA (Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures), http://www.fp7-federica.eu/

[96] N.M.M.K. Chowdhury, R. Boutaba, “A survey of network virtualization”, Computer Networks 54 (2010) pp. 862-876

[97] N.M.M.K. Chowdhury, M.R. Rahman, R. Boutaba, “Virtual Network Embedding with Coordinated Node and Link Mapping”, IEEE Infocom 2009

[98] Revealed: the environmental impact of Google searches, The Sunday times, January 11, 2009

[99] Alexander Wissner-Gross, How you can help reduce the footprint of the Web, The Sunday times, January 11, 2009

[100] W. Vereecken et al., “Power Consumption in Telecommunication Networks: Overview and Reduction Strategies”, IEEE Communications Magazine, June 2011

[101] Lantiq GmbH, Thomas Zettler, “Fundamental Energy Efficiency Measures for Home Gateways”.

[102] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, J.A. Maestro, "IEEE 802.3az: the road to energy efficient ethernet," IEEE Communications Magazine, vol.48, no.11, pp.50-56, November 2010.

[103] O. Goldschmidt, "ISP Backbone Traffic Inference Methods to Support Traffic Engineering", Internet Statistics and Metrics Analysis (ISMA) Workshop, San Diego, CA, December 2000.

[104] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot, “Traffic matrix estimation: existing techniques and new directions”, In Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM '02). ACM, New York, NY, USA, 161-174.

[105] Y. Zhang, M. Roughan, N. Duffield, A. Greenberg, “Fast accurate computation of large-scale IP traffic matrices from link loads”, In Proceedings of the 2003 ACM SIGMETRICS international conference on Measurement and modeling of computer systems(SIGMETRICS '03). ACM, New York, NY, USA, 206-217.

[106] J. Fang, Y. Vardi, Cun-Hui Zhang, “An Iterative Tomogravity Algorithm for the Estimation of Network Traffic”, Lecture Notes-Monograph Series Vol. 54, Complex Datasets and Inverse Problems: Tomography, Networks and Beyond (2007), pp. 12-23.

[107] Information and Communication Technologies: Enablers of a low-carbon economy, Alcatel Lucent white paper, [Online]: http://www.alcatel-lucent.com/eco/docs/CMO7526101103_ICT_Enablers-eco_EN_StraWhitePaper.pdf

[108] EU Code of Conduct on Energy Consumption of Broadband Communication Equipment, ”EU Stand-by Initiative”, URL: http://re.jrc.ec.europa.eu/energyefficiency/html/standby initiative broadband communication.htm

[109] EU Code of Conduct on Data Center Efficiency, URL: http://re.jrc.ec.europa.eu/energyefficiency/pdf/CoC%20data%20centres%20nov2008/CoC%20DC%20v%201.0%20FINAL.pdf

[110] D. Roche, ”Efficient broadband deployment & multi-service implementation”, Proc. Symp. on Transforming the Network through Global Collaboration (GSC13), Boston, USA, Jan. 2008.

[111] The Home Gateway Initiative (HGI), ”HGI puts Energy Saving on the Agenda”, White paper, http://www.homegatewayinitiative.org/ press/PR010.pdf.

[112] ITU Recommendation G.992.3: Asymmetric Digital Subscriber Line Transceivers 2 (ADSL2), http://www.itu.int/rec/T-REC-G.992.3/en, 2006.

[113] ITU Recommendation G.992.5: Asymmetric Digital Subscriber Line (ADSL) Transceivers, http://www.itu.int/rec/T-REC-G.992.5/en, 2006.

[114] G. Ginis, ”Low Power Modes for ADSL2 and ADSL2+,” SPAA021, Broadband Comm. Group, Texas Instruments, Jan. 2005.

[115] S.J.B. Yoo, "Energy Efficiency in the Future Internet: The Role of Optical Packet Switching and Optical-Label Switching," Selected Topics in Quantum Electronics, IEEE Journal of, vol.17, no.2, pp.406-418, March-April 2011.

[116] R.S. Tucker, "The Role of Optics and Electronics in High-Capacity Routers," Lightwave Technology, Journal of , vol.24, no.12, pp.4655-4673, Dec. 2006.

[117] R. Bolla, R. Bruschi, A. Cianfrani, M. Listanti, "Enabling backbone networks to sleep," Network, IEEE , vol.25, no.2, pp.26-31, March-April 2011.

[118] R. Tucker, “Will optical replace electronic packet switching?”, SPIE newsroom, 2007

[119] R. Bolla, R. Bruschi, F. Cucchietti, K. Christensen, F. Davoli, S. Singh, “The Potential Impact of Green Technologies in Next-Generation Wireline Networks: Is There Room for Energy Saving Optimization?”, IEEE Communications Magazine, vol. 49, no. 8, pp. 80-86, Aug. 2011.

[120] R. Bolla, R. Bruschi, A. Carrega, F. Davoli, “Green Technologies and the Art of Trading-off,” Proc. of the 1st Workshop on Green Communications and Networking (GCN 2011), co-located with IEEE Infocom, Apr. 2011, Shanghai, China.

[121] R.Bolla, R.Bruschi, A.Ranieri, “Green Support for PC-based Software Router: Performance Evaluation and Modeling,” Proc. IEEE ICC’09, Dresden, Germany, June 2009.

[122] R. Bolla, R.Bruschi, A. Carrega, F. Davoli, “An Analytical Model for Designing and Controlling New-Generation Green Devices”, Proc. of the 3rd Int. Workshop on Green Communications (GreenCom10) colocated with Globecom 2010, Miami, FL, USA.

[123] ITU-T Rec. G.993.5-2010, “Self-FEXT Cancellation (Vectoring) for Use with VDSL2 Transceivers.”

[124] V. Oksman, H. Schenk, A. Clausen, J.M. Cioffi, M. Mohseni, G. Ginis, C. Nuzman, J. Maes, M. Peeters, K. Fisher, P.-E. Eriksson, "The ITU-T's new g.vector standard proliferates 100 mb/s dsl," IEEE Communications Magazine, vol.48, no.10, pp.140-148, October 2010