

TREND Big Picture on Energy-Efficient Backbone Networks

Esther Le Rouzic, Raluca-Maria Indre
Orange Labs, Networks and Carriers, France
{esther.lerouzic ; raluca maria.indre}@orange.com

Luca Chiaraviglio, Francesco Musumeci, Achille Pattavina
Consorzio Nazionale Interuniversitario per le
Telecomunicazioni, Italy
luca.chiaraviglio@diet.uniroma1.it ;
{fmsumeci ; pattavina}@elet.polimi.it

Jorge López Vizcaíno, Yabin Ye
Huawei Technologies Duesseldorf GmbH, Germany
{Jorge.vizcaino ; Yeyabin}@huawei.com

Ward Van Heddeghem

Department of Information Technology of Ghent University
iMinds, Belgium, ward.vanheddeghem@intec.ugent.be

Andrea Bianco, Edoardo Bonetto, Michela Meo
Politecnico di Torino, Italy
{andrea.bianco ; edoardo.bonetto ; michela.meo}@polito.it

Felipe Jiménez
Telefónica Investigación y Desarrollo SA, Madrid, Spain,
felipej@tid.es

Filip Idzikowski
Technische Universität Berlin, TKN, Germany
filip.idzikowski@tu-berlin

Ruben Cuevas
Universidad Carlos III de Madrid, Spain
rcuevas@it.uc3m.es

Abstract— The ICT (Information and Communication Technology) sector has recently been identified as a growing contributor to worldwide greenhouse gases emissions and power consumption. This has triggered interest for more energy efficient ways to design and operate telecommunication networks. We present an overview of the solutions proposed within the European network of excellence TREND focusing on core and metro networks including data centers. Potential savings are presented and discussed with respect to their impact and applicability within the global picture.

Keywords— core networks ; data centers ; optical networks ; content distribution ; file distribution ; energy efficiency

I. INTRODUCTION

The impact of ICT on greenhouse gases emission has raised much interest over the last 5 years. First because it was suggested as a possible mean to help reducing the impact of other industries [1] and then also because its global footprint was shown to be non-negligible. In a recent study [2], the power consumption of ICT telecommunication networks is estimated to grow at rate of 10% per year, and its relative contribution to the total worldwide electricity consumption has increased from 1.3% in 2007 to 1.8% in 2012. In this context, researchers of the European network of excellence TREND [3] develop new, comprehensive, energy-aware approaches to networking including core, metro, access and home networks (wired and wireless) and data centers (Fig. 1).

In a previous paper [4] we focused on optical networks including access and core parts. In this paper we focus on the core and metro networks including data centers. Core and metro networks represent the inner backbone part of an operator telecommunication network. Their role consists of providing customers with connectivity to services of the operator (such as the Internet or its managed video services) or

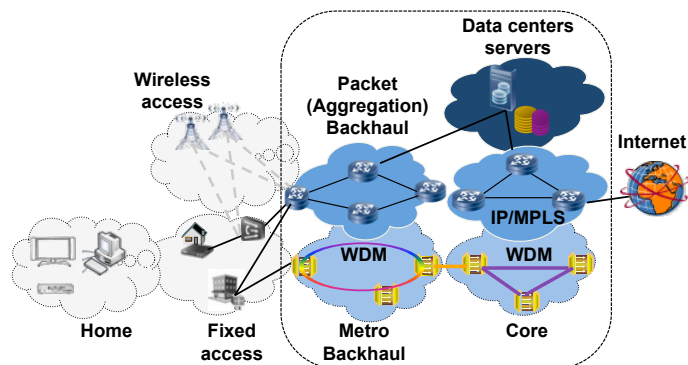


Fig. 1 Generic operator network architecture and domains studied within TREND. The frame represents the segment covered in this paper.

to services from other operators or over-the-top actors (such as content providers). It ranges from the first connection point up to the service platform or Peering Points (PPs) and includes transport equipment, optical transmission and data centers. We consider IP-over-WDM (Internet Protocol over Wavelength Division Multiplexing) networks, sub-wavelength photonic switching and Elastic Optical Networking (EON).

From the energy perspective, data centers have already been identified as a critical segment [5]. Meanwhile, core networks are expected to generate the biggest increase in electricity consumption [6] because their equipment number grows almost proportionally to the data traffic. Assuming similar exponential growth as observed over the last decade, the segment will also be critical on the 2020 horizon.

Our objective is to provide the TREND big picture on energy-efficient future backbone networks capable of sustaining traffic growth with reduced energy consumption. In Section II, we focus on more efficient architecture of networks and data centers. In Section III, we present energy efficient ways to operate the networks. We conclude the paper with a

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257740 (Network of Excellence “TREND”).

discussion and a wrap-up of the potential savings (Section IV). The list of references consists mainly of TREND publications, which in turn contain further related work.

II. MORE EFFICIENT ARCHITECTURES AND DESIGN

A. Energy-Efficient Network Design

1) Optical bypass to save transit in IP routers

In typical IP-over-WDM core networks, the IP nodes represent the largest share in the total power consumption. It has been shown to be as high as 60% of the total power consumption [7].

Therefore, any solution that allows reducing the average number of IP hops required to route a demand will usually decrease the power consumption of the core network. Optically bypassing the IP routers is a well-known technique to reduce the IP hop count. The idea is that traffic not intended for the IP node, remains in the optical layer and thus bypasses the IP. This allows us to reduce the capacity of the router and the associated power consumption.

For sufficiently large demands, this technique was shown in [7] to allow energy savings of up to 50% in the pan-European Géant network. A more extensive analysis in [8] shows that these savings are for a great deal depending on the mesh degree of the network, the ratio of the average demand over the line rate, and the average physical link length. High(est) savings are achieved for the combination of demands close to (but below) the channel line rate, sparsely meshed networks (corresponding to realistic networks with node degrees of around 3), and not very long link lengths.

2) Energy awareness in the design of optical core networks

The savings obtained by optical bypass suggest that more efficient design of networks may be obtained when including energy as a metric. Several design choices are available to improve the energy efficiency of the core network segment [9]. Identifying the most important design parameters is thus required.

The most important choice is the design objective. Setting minimization of energy consumption as the design objective brings of course a significant reduction of energy consumption with respect to choosing performance optimization (like traffic load minimization) as the design objective. Achievable energy savings range from 10% to 30% depending on the network topology and the traffic load [9].

Another design choice concerns the routing of the traffic demands: single-path vs. multi-path routing. Adopting multi-path routing requires that each traffic demand is split into several sub-demands routed independently, potentially following different paths from the source to the target node. Thus, multi-path routing can achieve a more effective traffic grooming, resulting in fewer network interfaces and as a consequence in lower network energy consumption. Savings are significant only for low traffic load; otherwise the savings reach roughly 10% [9].

In summary, the energy consumption of core networks can be significantly reduced if appropriate choices are performed during the design phase. However, some of the energy wise

choices may reduce network performance. For example, optical bypass may reduce network availability because in case of a failure, it prevents from re-routing at the intermediate router that has been bypassed; thus a trade-off between energy reduction and performance should be sought for.

3) Optical aggregation with sub-wavelength photonic switching technologies

Using optical switching of lightpath instead of electronic at intermediate network nodes eliminates the need for optical to electronic to optical (O/E/O) conversion in the core nodes of the network. On the downside, this suffers from both scalability and efficiency issues especially in metro and backhaul networks, where there may be higher discrepancy between the effective traffic and the wavelength channel granularity. Aggregation is then needed to avoid dedicating a wavelength to a single source-target pair leading to severe bandwidth waste. In current networks the bandwidth of every wavelength channel is dynamically shared between several source-target pairs thanks to electronic switching. Sub-wavelength photonic switching technologies are also providing this capability.

In its simplest form, sub-wavelength switching consists of dividing the capacity of each wavelength channel into synchronous fixed-size time-slots. Time-slotted WDM architectures have been demonstrated as a viable approach to provide all-to-all connectivity in ring-based metro optical networks, e.g., [10]. In such networks, slots propagate on the ring and nodes can communicate by adding/dropping data to/from in-transit slots. By allowing intermediate nodes to be bypassed in the optical domain, the time-slotted switching architecture has the potential of reducing the energy consumption of today's electronic-based metro networks by up to 75% [11]. However, the time-slotted solution requires slot synchronization of all network nodes, which may be difficult to achieve in a meshed Wide Area Network (WAN).

Optical Burst Switching (OBS) with reduced processing overhead and switching constraints may be an interesting solution for meshed networks. In OBS networks, incoming IP packets are aggregated into larger optical bursts at the edge of the OBS network. Unlike time-slotted solution burst of data may have variable length and can be sent asynchronously. Several flavors of control exist and we showed that the most promising solutions are the ones that avoid burst loss [12]. By eliminating O/E/O conversion at intermediate nodes, OBS is likely to considerably reduce power consumption. A recent proposal for sub-wavelength switching at burst granularity has shown by means of a detailed comparative study that sub-wavelength switching solutions are expected to reduce the energy consumption by some 50% when compared to today's electronic-based WANs [13].

B. Energy Efficient Network Architecture

1) Energy-Efficient Content Distribution

Traffic volume generated by multimedia content represents the major part of the traffic and this part is expected to grow in the future. As an example, Cisco estimates that the sum of all forms of videos will represent 86% of the global consumer traffic by 2016 [14]. Content distribution is thus an important

issue for the energy consumption. Thus, an energy-efficient design of content distribution architectures managed by an Internet Service Provider (ISP) has been considered in [15]. The aim is to find the best locations in the ISP topology where to place the content requested by users from the Internet Peering Points (PPs) in order to minimize the energy consumption of the ISP infrastructure including transport, storage caches and servers.

An example of the assumed scenario is reported in Fig. 2. The topology is composed of different levels, being level 0 the PP providing connectivity with the rest of the world and the last level the set of users requesting content. Without content distribution, the content is sent directly from the PP to the users (Fig. 2-(a)). This procedure has to be performed each time the content is requested by users (at time t_1 and t_2 in the figure), since the content is stored outside the network managed by the ISP. On the contrary, with the proposed approach, the content is first fetched into the caches at level j (Fig. 2-(b)), stored for a certain amount of time and then accessed by users when needed (Fig. 2-(c)). In this way, a decrease of traffic is possible in the portion of the network between the first level and level j , which in turn triggers a reduction of the energy consumption.

The decision about where to store the content is taken on the basis of the energy consumption and the content popularity. In fact, very popular content is likely to be placed close to users to minimize the cost of transporting information inside the network. However, a tradeoff exists since the more the content is placed close to the users, the higher number of replicas are required. The preliminary results obtained over a realistic case study show that energy savings of more than 8% are achievable. These savings are obtained on the entire network of the ISP, ranging from the core to the access. Consequently, huge monetary savings are possible, corresponding to more than 750 k€ of saved money in one year. At the same time, a reduction of 18% in the bandwidth required at the PP is possible.

2) Energy-optimal file distribution

Recent studies [6] demonstrate that homes and organizations (i.e., end-hosts) are responsible for 75% of the overall Internet energy consumption. Furthermore, existing file distribution services such as peer-to-peer (P2P) file sharing applications, one-click-hosting services, software releases, etc., represent a major fraction of the current Internet traffic [16]. Since file-sharing applications are typically run by end-host devices (e.g., PCs or laptops), the combined effect of the two previous arguments suggests that file-sharing applications are

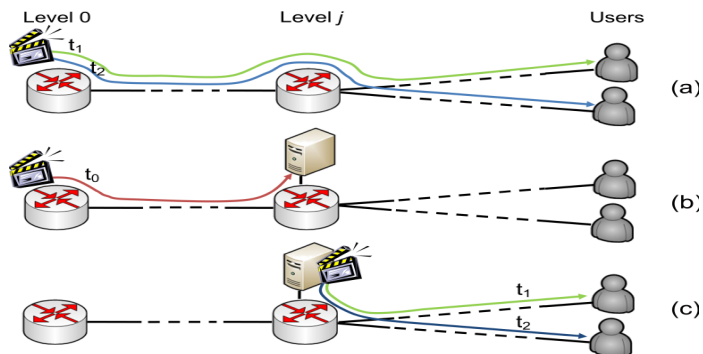


Fig. 2 Traditional content distribution (a) and content caching approach (b) (c).

responsible for a significant fraction of the overall energy consumption in the Internet.

The majority of previous studies in the area of optimizing file distribution services have mainly focused on minimizing the download time [17]. However, those algorithms, designed to minimize the download time, are not optimal in terms of energy consumption.

Performance of energy-efficient file distribution algorithms was analyzed in [18]. In particular we proposed to use the upload capacity of hosts, which are on for downloading a block of the file to serve blocks to other hosts in a P2P fashion.

The obtained results show that the proposed algorithms reduce the energy consumption in a file distribution process with respect to any centralized file distribution schemes (which are widely used). In particular, the simulations show that, even in scenarios for which our algorithms were not designed for (e.g., considering heterogeneous energy consumption or network congestion), the proposed collaborative P2P schemes achieve significant energy savings with respect to largely used centralized file distribution systems. These savings range between 50% and two orders of magnitude, depending on the centralized scheme under consideration.

C. Energy efficient Data Centers

1) Reducing the Power Consumption of Network Devices in Large Data Centers

Recent studies have shown that between 10% and 20% of power in a data center is consumed by network devices. Thus, it makes sense to investigate solutions reducing the power consumption of these devices. To this end, we proposed a solution, called VMPlanner [19], to efficiently optimize the network power consumption in a cloud data center. VMPlanner performs two main tasks: the virtual machines placement inside the data center, and the traffic flow routing between the virtual machines. At the end of this procedure, as many network elements as possible are switched off while guaranteeing Quality of Service (QoS) constraints.

The problem is formulated as an Integer Linear Problem and solved sub-optimally by decoupling in three different steps. In particular, first the virtual machines are grouped into sets that exchange a minimum amount of traffic. Then, as the second step, the virtual machines are mapped to the server infrastructure in order to minimize the traffic that is exchanged between the servers. Finally, the traffic generated by the virtual machines is routed in such a way that as many network elements as possible can be turned off. Results demonstrate the efficiency and efficacy of the proposed solution, with up to 60% of saving achievable compared to a solution in which all the network devices are always powered on [19].

2) Self-Organizing Algorithm for Energy Saving in Data Centers

The ever increasing demand for computing resources has led companies and resource providers to build large warehouse-sized data centers, which require a significant amount of power to be operated: 120 billion of kWh in the USA alone, as stated in the report published by the USA Environmental Protection Agency [20].

Many efforts have been devoted to the efficiency improvement of the physical infrastructure of the data centers. Indeed, only a fraction of the power that enters the data centers is actually used to support the workload, while a consistent amount of power is used by the facilities for power supply and distribution, and by the cooling infrastructure. However, the physical efficiency is not enough to guarantee the overall efficiency of the data center that has to be also related to the computational efficiency that, in its turn, depends on the workload distribution over the servers. A given load evenly distributed over a large number of servers that work at low utilization requires more energy than the same load groomed over a smaller number of servers that work at higher utilization. This observation, combined with the possibilities offered by virtualization, has led to many approaches to consolidate load on as few servers as possible. The key issues are then the strategies to properly allocate virtual machines to servers and the methods to face the possibility that servers, which work at high utilization, become overloaded due to fluctuations of the amount of resources requested by the running virtual machines.

The solution proposed in [21] consists of delegating part of the decision processes to the servers, so that the allocation and migration strategies are implemented in a distributed fashion with the contribution of all the servers. The servers use only the local knowledge about their own load and make simple probabilistic decisions about their availability to host a new virtual machine or the need to migrate some virtual machine to other servers. The probability of accepting a new VM or migrating one depends on the server load, so that overload conditions can be avoided. By acting this way, the system, as a whole, naturally scales and adapts to the varying conditions. The achieved energy saving depends on the workload profile varying over time and, thus, on the degree of server consolidation that can be achieved. Average daily saving is typically between 30 and 40%.

III. ENERGY EFFICIENT NETWORK OPERATION

A. Dynamic adaptation

Traffic varies over time giving the opportunity for network devices to save energy by entering sleep modes or reducing rate of operation during low demands hours.

1) Adaptive routing at IP and WDM layers

Network topologies in the core usually form a mesh, where more than one path is present between a node pair. Dynamic choice of the path for the traffic demands arriving at the network can be used to increase the number of idle or lightly loaded devices during the low demand hours [22]. It is difficult to switch off complete nodes of core networks, because there is usually (aggregated) traffic originating from or targeted to them. However, traffic can be rerouted from some links, and delivered to the target nodes via alternative paths. Furthermore, many IP links are composed of bundled links of the lower layer, so it is possible to temporarily decrease their capacity and deactivate the excessive resources. Furthermore, the rate of both the link and node equipment could be reduced to decrease its power consumption.

We presented a wide set of adaptive routing solutions working in both the IP layer and the optical layer in [23]. An

evaluation study based on the France Telecom network scenario and medium traffic forecast for the year 2020 has shown that the different solutions utilizing adaptive routing in the IP layer achieve comparable savings. 39-45% of energy consumed by line cards over a working day can be saved with respect to the static base network. We made sure that the over-provisioning targeted by the adaptive solutions remains at the same level as in the static base network dimensioned for the peak traffic. Furthermore, we compared these savings with the savings of a simple solution called Fixed Upper Fixed Lower (FUFL), where no dynamic routing is used, but only dynamics of traffic is utilized. FUFL achieves energy savings of 15% in the considered scenario. The smaller savings in comparison to the results obtained by the more sophisticated solutions are balanced by the simplicity of application of this distributed solution, as it requires only local knowledge to make decisions about activation or deactivation of resources.

The percentage of energy savings achieved by the solutions in the optical layer (with respect to all devices targeted to be put into sleep mode) is similar or even higher. However, the absolute values of the savings in kWh are much smaller than those in the IP layer, as optical line amplifiers consume significantly less power than the router line cards [7].

2) Improved spectral efficiency of new coherent formats at the WDM layer

The increase of traffic experienced and forecasted in core networks has resulted in a strong research on new technologies aimed at increasing the spectral efficiency of the transmission networks in order to increase fiber capacity. These technologies are based on the use of several key blocks among which the following may have an important impact on energy efficiency:

- Coherent reception with DSP (Digital Signal Processor) assisted impairment compensation.
- Advanced modulation formats with higher numbers of transmitted bits per symbol.
- Multicarrier structures with improved spectrum efficiency (super channel or Orthogonal Frequency Division Multiplexing OFDM).

Indeed the generalization of signal processing with the DSP on the one hand requires additional power consumption compared to previous transmission generation that did not use it. On the other hand the improved spectral efficiency and electronics, may instead improve the global energy consumption at the WDM transmission layer due to better power efficiency (Watt/Gb/s).

Currently available 100 Gb/s solutions based on coherent implementations will be a commonplace in the very short term. Also solutions offering up to 200 Gb/s capacity on a single 50 GHz frequency slot are starting to be deployed, thus offering 20 times more capacity than traditional 10 Gb/s technologies. Flexible grid allows increasing the spectral efficiency even to higher figures.

For coherent technologies, the DSP is a major power contributor and has a direct relation with the signal symbol rate: higher order modulation format translates in a better power efficiency. This could be exploited using the concept of

“software defined” transponders that along with flexible grid networking form the basis of the EON.

3) *Elastic Optical Networking*

Recently, the concept of EON, based on OFDM, for example, and coherent detection, has been introduced as a promising candidate for the operation of future optical transport networks. EON can bring significant advantages in terms of energy efficiency thanks to the variable lightpath capacity and the adaptive modulation formats. The finer granularity allows a better adjustment of the allocated capacity by expanding or contracting the channel bandwidth according to the actual user demand, whereas the adaptive modulation enables the choice of modulation format according to the demand and distance, thus minimizing the number of regenerations in the network. Furthermore, EON also improves the spectral efficiency and reduces the network blocking with respect to fixed-grid scenarios, which has also an impact on the energy efficiency of the network.

Studies in terms of energy efficiency for EON have been carried out for different network scenarios and conditions in [24] (i.e., different-sized network topologies, static and dynamic traffic operations, unprotected and protected networks, and different traffic loads). One of the main conclusions from the different studies is that EON can considerably benefit in terms of power efficiency when exploiting the traffic variations in the network. For instance, EON can specially benefit from novel protection scheme, proposed in [25], where the rate of the backup transponders is adapted to the current required bandwidth requirements to reduce the energy consumption. Power consumption reductions up to 18% can be achieved by this scheme with respect to the conventional dedicated protection 1+1 (DP 1+1).

B. *Energy-efficient protection schemes*

1) *Quality of protection*

In long-haul optical networks, the most common and secure strategy to provide resilience is implementation of a DP 1+1, scheme, where the data is duplicated and transmitted on two link-disjoint paths. This scheme requires the reservation of twice the spectral resources for working and protection paths, and the deployment of redundant transponders, which are simultaneously active and consuming power. However, in many cases clients may not require such a high level of reliability for their service. Accordingly, the heterogeneity of protection requirements requested by clients could be exploited to enhance the energy efficiency of the network with respect to the conventional DP 1+1. A differentiated quality of protection (Diff QoP) was proposed in [25] and evaluated for both fixed-grid WDM and flexible grid OFDM-based EON. The provisioning of Diff QoP is based on the definition of different QoP classes corresponding to different protection levels: DP 1+1, DP 1:1, shared protection and best-effort protection.

In the study, different traffic load conditions with different percentages of QoP traffic classes have been evaluated corresponding to the predominance of different types of client: i) big corporations, ii) small- and medium-enterprises, and iii) intermediate scenario. Energy savings up to 21% can be achieved with respect to DP 1+1 scheme. Even though power savings can be obtained for all the different traffic conditions,

the degree of power savings strongly depends on the transmission technology, total traffic, and distribution of traffic classes. Savings are more significant at high traffic load, and for scenarios with lower percentage of client demands requiring maximum protection. More details can be found in [25].

2) *Protection based on sleep-modes*

In this context, considering that protection resources (i.e., WDM transponders, IP routers line cards, etc.) can be set into a low-power sleep-mode, high energy benefits can be obtained. In [26], a comparison between different protection strategies has been carried out from the power consumption point of view, by accomplishing a power-minimized network design. Dedicated vs. shared as well as link vs. path protection strategies have been compared, developing for each of the four scenarios an Integer Linear Programming based formulation aiming at minimizing the network power consumption. It has been found in [26] that up to about 60% of power savings, according to the protection strategy and traffic scenario, can be obtained by setting protection devices into sleep-mode, especially thanks to the possibility of saving large amounts of the power spent in IP routers. Moreover, it has been also found that, by employing sleep-mode for protection devices, it is possible to guarantee network resilience for a small (1–2%) additional power expenditure compared to non-resilient scenarios. Further results and discussion can be found in [26].

C. *Cooperation*

1) *Content Distribution and Network Management for Energy-Efficiency*

Internet Service Providers are becoming sensitive to reduce their power consumption due to the electrical power consumption of the network equipment. At the same time, also Content Providers (CPs) are experiencing huge energy costs due to the massive deployments of servers and computational machines. To this purpose, we propose a new solution in which ISPs and CPs cooperate together to reduce jointly the power consumption. The ISPs manages a set of routers and links, while the CPs can operate on servers. By properly turning on/off network and server equipment, it is possible to reduce the overall power consumption.

Traditionally, ISPs and the CPs are not willing to share sensitive information such as the network topology and the servers' load. Therefore, the proposed approach is distributed over the ISP and the CP to limit the amount of exchanged information.

We evaluate our model with realistic ISP topologies and power figures. Simulations show that large power savings are achievable, up to 71% relative to a classic formulation of the joint problem, which minimizes only the users' delay. Moreover, we have shown that a preferential server placement saves an additional 15% over the overall power consumption with respect to a random one. Finally, our results indicate that the distributed solutions are near-optimal for all considered scenarios, with a maximum power saving of 54% with respect to the classic formulation of the joint problem. More details can be found in [27].

IV. CONCLUSION

The schemes proposed in this paper could lead to the (estimated) energy consumption reductions¹ summarized in Fig. 3. These numbers are provided only as purely indicative ones (according to corresponding references), since their relative impact on the global picture depends on the consumption of the segment they address and the considered network scenario. This is part of a future work.

Recommendations and next steps

Cost savings generated by energy consumption reduction is already an incentive for network operators. Regulation actors clearly play an important role with a positive impact on energy when asking vendors and operators to meet certain standards or negative impact on the total network energy consumption, for instance when it leads to the deployment of multiple infrastructures or equipment to avoid a dominant position of an operator.

Energy consumption reduction should be though beyond the actors own equipment keeping in mind that savings on one part of the networks may result in an increase in other parts. End-to-end vision is inevitable to make the most of the solutions. It is also required to maintain high levels of QoS. Core networks are especially critical segments where it is not possible to trade QoS for energy savings. As a result the impact of any proposed solution should be evaluated with respect to its possible impact on failure rates of network devices.

Taking a global view on energy-efficient optical networks is a complex task and remains an open issue for further investigations.

V. REFERENCES

- [1] GeSI and BCG, "SMARTer2020: The Role of ICT in Driving a Sustainable Future," 2012. <http://www.theclimategroup.org>.
- [2] S. Lambert et al., "Worldwide electricity consumption of communication networks," *Optics Express*, vol. 20, no. 26, pp. B513–B524, Dec. 2012.
- [3] TREND web site: "<http://www.fp7-trend.eu>".
- [4] E. Le Rouzic et al., "TREND towards more energy-efficient optical networks," in *Proc. ONDM 2013*.
- [5] J. Koomey, "Growth in Data center electricity use 2005 to 2010," Oakland, CA: Analytics Press, August 2011. <http://www.analyticspress.com/datacenters.html>
- [6] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch, "Energy consumption of telecommunication networks and related improvement options," *IEEE J. Sel. Top. in Quantum Electronics*, vol. 17, no. 2, pp. 285–295, 2011.
- [7] W. Van Heddeghem et al., "Power consumption modeling in optical multilayer networks," *Phot. Net. Comm.*, vol. 24, no. 2, pp. 86–102, Oct. 2012.
- [8] W. Van Heddeghem et al., "Power Consumption Evaluation of Circuit-Switched Versus Packet-Switched Optical Backbone Networks," submitted to *OnlineGreenComm 2013*.
- [9] A. Bianco, E. Bonetto, and A. Ahmad, "Energy awareness in the design of optical core networks," in *Proc. OFC 2013*.
- [10] D. Chiaroni et al., "Packet OADMs for the next generation of ring networks," *Bell Lab. Tech. J.*, vol. 14, no. 4, pp. 265–283, Feb. 2010.
- [11] A. Bianco, T. Bonald, D. Cuda, and R.-M. Indre, "Cost, power consumption and performance evaluation of metro networks," *J. of*

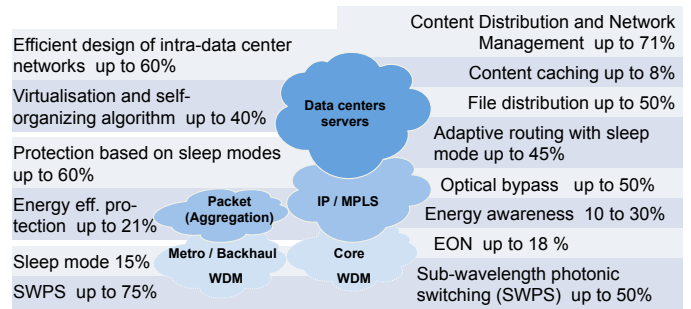


Fig. 3 Summary of the estimated energy reductions¹ presented in the paper.

- [12] E. Bonetto, A. Triki, E. Le Rouzic, B. Arzur, and P. Gavignet, "Circuit switching and time-domain optical sub-wavelength switching technologies: Evaluations on the power consumption," in *Proc. SoftCOM 2012*.
- [13] D. Cuda, R.-M. Indre, E. Le Rouzic, and J. Roberts, "Building a low-energy transparent optical wide area network with "multipaths"," *J. of Optical Communications and Networking*, vol. 5, no. 1, pp. 56–67, Jan. 2013.
- [14] "Cisco visual networking index: Forecast and methodology, 2011–2016," 2012.
- [15] R. Modrzejewski et al., "Energy efficient content distribution in an ISP network," accepted for *Globecom 2013*.
- [16] C. Gkantsidis, T. Karagiannis, and M. Vojnović, "Planet scale software updates," in *ACM SIGCOMM*, 2006.
- [17] T. Langner, C. Schindelhauer, and A. Souza, "Optimal file-distribution in heterogeneous and asymmetric storage networks," in *Proc. SOFSEM 2011*.
- [18] K. Verma, G. Rizzo, A. Fernández Anta, R. Cuevas and A. Azcorra, "Greening the internet: energy-optimal file distribution," in *Proc. NCA 2012*.
- [19] W. Fang, X. Liang, S. Li, L. Chiaraviglio, and N. Xiong, "VMPlanner: optimizing virtual machine placement and traffic flow routing to reduce network power costs in cloud data centers," *J. Computer Network*, vol. 57, no. 1, pp. 179–196, Jan. 2013.
- [20] U.S. Environmental Protection Agency, "Report to Congress on Server and Data Center Energy Efficiency," Public Law 109-431, Aug. 2007. http://www.energystar.gov/ia/partners/prod_development/downloads/EP_A_Datacenter_Report_Congress_Final1.pdf
- [21] C. Mastroianni, M. Mèo, and G. Papuzzo, "Analysis of a self-organizing algorithm for energy saving in data centers," in *Proc. HPPAC*, 2013.
- [22] F. Idzikowski, S. Orłowski, C. Raack, H. Woesner, and A. Wolisz, "Dynamic routing at different layers in IP-over-WDM networks – Maximizing energy savings," *Optical Switching and Networking, Special Issue on Green Comm.*, vol. 8, no. 3, pp. 181–200, July 2011.
- [23] F. Idzikowski et al., "TREND in Energy-Aware Adaptive Routing Solutions," submitted to *IEEE Communications Magazine*, Aug. 2012.
- [24] J. López Vizcaíno et al., *Green Networking and Communications: ICT for Sustainability*, chapter on "Energy efficiency improvement with the innovative flexible-grid optical transport network," Auerbach Publications, CRC Press, Taylor & Francis Group, USA, in press.
- [25] J. López Vizcaíno et al., "Differentiated quality of protection to improve energy efficiency of survivable optical transport networks," in *Proc. OFC 2013*.
- [26] F. Musumeci, M. Tornatore, J. López Vizcaíno, Y. Ye, and A. Pattavina, "Energy-efficiency of protected IP-over-WDM networks with sleep-mode devices," in *J. of High Speed Networks*, vol. 19, no. 1, pp. 19–32, 2013.
- [27] L. Chiaraviglio and I. Matta, "Energy-Aware Network Management and Content Distribution," in J. Wu, S. Rangan, and H. Zhang (Editor), *Green Communications: Theoretical Fundamentals, Algorithms, and Applications*, CRC Press, 2012.

¹ The reduction achieved by Content Distribution and Network Management refers to power consumption, and the reduction achieved by file distribution refers to energy per bit.