Power Savings versus Network Performance in Dynamically Provisioned WDM Networks

Pawel Wiatr, Paolo Monti, and Lena Wosinska, Royal Institute of Technology KTH

ABSTRACT

The role of the ICT sector in our daily life is causing a significant growth of the power consumed by the network equipment. In this context, transparent WDM networks represent a promising solution for reducing the power consumption of telecom networks. For this reason transparent WDM networks attract a lot of interest and several power-aware routing and wavelength assignment algorithms are available in the literature. These approaches, however, seem to consider power minimization as the only parameter to be optimized. Little or no attention is given to other important network aspects, e.g., connection requests blocking. Such approaches may not always be advisable, especially in core WDM networks where high performance levels must be considered in the first place. This article aims at providing a different insight to the PA-routing problem. It follows the intuition that in some cases relaxing the power minimization constraint can have beneficial effects on the overall network performance, especially on the blocking probability. A novel approach, referred to as Weighted Power-Aware Lightpath Routing (WPA-LR), is proposed and evaluated using two continental core networks. Our results confirm the presence of a trade-off between power savings and blocking probability.

INTRODUCTION

Historically, power efficiency has always been a critical issue where energy storage capabilities are limited (e.g., in wireless communications). In the last few years, the interest of the research community in more energy-efficient, or *green*, strategies has broadened to also include wired networks. This trend is mainly motivated by the role and the rate of penetration of the information and communication technology (ICT) sector in our daily life. Consequently, network equipment is becoming a significant contributor to the global energy consumption making the power efficiency of communication networks increasingly important. Power reduction in telecom networks can be addressed by: development and deployment of power efficient equipment and transport technologies as well as power aware control, management and provisioning of connection requests within the network.

It is widely recognized that optical networks offer power efficient transport technology and play an important role in reducing the power consumption of telecom networks [1]. In fact, replacing conventional electronic technology with their optical counterpart offers noticeable reductions in the overall network power consumption [2]. The main benefit of transparent optical networks derives from their ability to avoid optical-electrical-optical (O/E/O) conversion. Their power efficiency can be further increased by a proper network design [3–5] and green routing approaches [6–8] where the energy saving is based on minimization of the number of active network elements that need to be powered on in order to guarantee the required connectivity. According to these methods power reduction can be made possible by putting unused network elements in a power saving, or *sleep*, mode of operation. Although not available yet in most network devices, support of sleep mode is advocated by current efforts from standardization bodies, governmental programs, and research papers that recommend the use of sleep mode in the different network devices [9, 10].

There are currently many proposals to apply these solutions in both static and dynamic provisioning scenarios. In the static paradigm [3–5], these methods are used to provision traffic in already deployed network to minimize the number of active elements, thus lowering the overall power consumption. On the other hand, in a dynamic scenario [6–8] a selective sleep mode option for some network elements (usually links with traffic below a certain threshold) is considered while the network is in operation. This requires real time re-routing of the traffic traversing the network elements that are in sleep mode, without disturbing the connection requests already provisioned.

Regardless of the provisioning scenario, the focus of the energy efficient methods mentioned above is put on power minimization only. Little or no attention is given to the impact that green solutions have on other network performance metrics. However, ignoring the possible impact

0163-6804/12/$25.00 © 2012 IEEE

of power minimization on other network performance parameters may not always be acceptable.

The trade-off between energy saving and network blocking performance can be illustrated by the following example. Consider the scenario presented in Fig. 1, where a certain number of lightpaths needs to be dynamically provisioned in a WDM network with seven nodes and eight bidirectional fiber links. In the example, the term lightpath refers to a transparent (optical) connection between a pair of network nodes. It is assumed to have one fiber link in each direction, with two wavelengths per fiber. Moreover, for the sake of simplicity it is considered that all fiber links have the same length and the power consumption of a fiber link (proportional to its length) is the same, while nodes, once active, consume the same fixed amount of power.

First, consider a provisioning strategy where the only objective is power minimization, i.e., among the various candidate paths, the one with the lowest impact on network power consumption will be selected. At the beginning the network is assumed to be empty and all devices are not in use. When connection request R1 between node E and node C arrives, the path at minimum power cost, E-D-C using two fiber links and one node, is reserved. Then connection request R2 between nodes A and G arrives. In this case the path at minimum power cost is A-E-D-C-G, taking advantage of the fact that fiber links E-D and D-C are already in use. Assume now a third connection request R3 between nodes B and G. With the same rationale the provisioned path will traverse B-C-G. However, should a fourth connection request R4 between nodes C and G arrive, it would be blocked, since both C-D and C-G wouldn't have wavelength resources available. This provisioning strategy is able to save 35 percent (3 of 8 links are not used) of power consumed by fiber links, 14 percent (1 of 7 nodes is not used) of the power consumed by the network nodes, at the cost of having one out of four connection requests blocked (i.e. 25 percent blocking). If, on the other hand, a provisioning strategy based solely on the minimum number of hops is considered, the outcome is the following. R1 is routed along path E-D-C. R2 is assigned path A-F-G, while R3 is routed along B-C-G. As a result path C-G is still available for R4. With this provisioning strategy the potential power saving is lower since only 12 percent of links (1 of 8 links) are not in use and all nodes need to be powered. The blocking performance, on the other hand, is better in this case, because the minimum hop routing is able to avoid the bottlenecks created on fiber links C-D and C-G by the strategy focusing on energy saving.

This simple example shows that power minimization has an adverse impact on the length of the provisioned lightpaths since it forces the selection of already used links in order to avoid powering on the unused ones. In other words, in an attempt to reduce power consumption the allocated paths become on average longer, inadvertently creating bottlenecks in the network. This is in contradiction with the goal of classical (i.e., not power aware) routing and wavelength assignment (RWA) algorithms that tend to minimize the resource usage within the network in order to minimize the blocking probability.

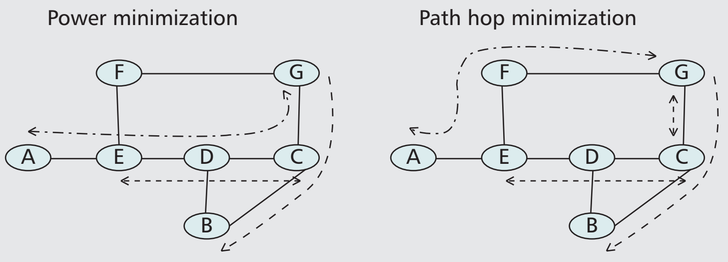

Requests		Calculated routes	
No	Source - dest.	Power minimization	Path hop minimization
R1:	E - C	E-D-C	E-D-C
R2:	A - G	A-E-D-C-G	A-E-F-G
R3:	B - G	B-C-G	B-C-G
R4:	C - G	Blocked	C-G

Figure 1. *Power vs. blocking, a trade-off example.*

This article addresses this issue by investigating the impact of power minimization on the network performance and proposes a non-conventional solution to the power aware routing problem. This approach, referred to as Weighted Power-Aware Lightpath Routing, is based on the intuition that relaxing the power minimization constraint can have beneficial effects on the overall network performance by reducing resource fragmentation in the network and, consequently, lowering the blocking probability. The proposed algorithm leverages on a cost function that weights the power status of network elements versus the wavelength usage. The algorithm has been tested on two core networks, the Pan-European core network (i.e., COST 239 [11]) and the NSFNET topology [12]. Performance results reveal the presence of a trade-off between power savings and blocking probability. It is also shown that a "binary" approach based only on the information about the power status of a network elements (powered on vs. sleep mode) might result in unsatisfactory network blocking.

The article is organized as follows. We provide an overview of the current research efforts in the field of green connection provisioning. We describe the proposed power aware strategy whose performance results are analyzed. Finally, we present some concluding remarks.

OVERVIEW OF GREEN CONNECTION PROVISIONING STRATEGIES

In the recent years a significant effort has been invested into studying the energy efficiency of communication networks. The concept of "greening the Internet" is introduced in [13]. An energy-aware routing scheme for a static case is proposed in [14] where the power consumption of a router is modeled by considering the energy consumption of chassis, the number of line cards in the chassis and the number of ports in a line card. The total power consumption of the router

Ignoring the possible impact of power minimization on other network performance parameters may not always be acceptable.

fabric is then minimized by maximizing power efficiency of each one of these parts. The work in [15] introduces the concept of power profile routing based on generic models for the power consumption of the routers, while in [16] the benefits of load adaptive networks are described. According to these concepts, if resource allocation can be done resorting to some power-consumption profiles that are a function of the traffic, the network power efficiency can be significantly improved. The authors in [17] propose a different approach to reduce the power consumption where idle line cards and their corresponding lightpaths are powered off at a low traffic load. The presented results show that rerouting demands in the IP layer has the strongest impact on the energy savings. Allowing reconfiguration in the optical domain barely brings any extra benefit in the scenarios analyzed in [17]. The results also indicate that energy should be included in the objective function whenever IP routing weights are reconfigured in the operational phase.

In the context of optical networks, the work in [2] evaluates the energy savings achievable through the use of wavelength routed optical transport networks in a realistic framework of the network of the future. One important result that is found is the possibility to considerably reduce power consumption by relying on optical transmission technologies. For this reason, in order to make further improvements, energy efficiency in the optical layer has attracted a lot of attention and a wide range of topics are addressed in the literature.

Some of the work on energy-aware routing focuses on the problem of determining the resources that are under-utilized and that can accordingly be put in sleep mode to save energy. The concept of reducing power consumption in a static provisioning case by minimizing the number of powered on network elements is addressed in [3] by using a mixed integer linear program (MILP) formulation. The results show that design criteria aimed at power saving can be very effective. The work in [7] proposes an approach by which clusters of wavelength routing nodes in an optical backbone network adopt a sleep mode. While sleeping, a node is still able to support the previously assigned pass-through traffic, as well as to transmit and to receive the lightpaths originating and terminating at the node. Due to the proposed sleep cycles, the connectivity of the network decreases and the probability of requests being blocked increases. For this reasons the authors investigate the trade-off between connection request blocking and energy efficiency in an anycast provisioning scenario. The work in [6] presents a distributed strategy that can be used in WDM networks to select the fiber links to be set in sleep mode and to reroute the lightpaths accordingly. In the same work GMPLS extensions for advertising the fiber links in sleep mode are also introduced. The authors show that considerable power savings are achieved at low loads, without affecting the network performance in terms of blocking. The use of sleep mode for protection resources was introduced in [4] where the authors address the energy-efficient and survivable network design problem with dedicated path protection. Their results show that power savings (up to 25 percent) are achievable by properly provisioning lightpaths in the network so that devices, such as Erbium doped fiber amplifiers (EDFAs), supporting only protection lightpaths can be set in sleep mode. The same intuition is used in a dynamic provisioning scenario leading to equally promising power savings results [8]. The authors in [5] on the other hand consider an energy-efficient network design problem with shared path protection.

The use of sleep mode is indeed controversial, given that the technology to enable this feature is not yet mature, and that this technique might be seen as problematic by network operators due to the time necessary to put to sleep and to wake up a network device.

None the less the studies presented so far evaluate the maximum energy savings that could be achieved by leveraging on the sleep mode operation. A comprehensive survey on the topic of green telecom networks can be found in [18].

WEIGHTED POWER-AWARE LIGHTPATH ROUTING STRATEGY

This section introduces a power aware routing strategy that is used to study the impact that power minimization has on the other network metrics. Before going into the details of the proposed routing strategy, the power model used in the study is presented.

POWER MODEL

This study assumes a transparent WDM network where nodes are equipped with optical cross-connects (OXCs) and fiber links comprise a number of EDFAs, depending on the fiber link length. Connection requests are provisioned all-optically from source to destination (i.e., without optical-electronic-optical [O-E-O] conversion at intermediate nodes). Based on these assumptions, the power required to provision one connection request is equal to the sum of: the power consumption of the transceiver, the power needed to optically switch the signal at intermediate nodes, (i.e., OXCs), and the power consumed by the EDFAs along each fiber link in the path. The overall network power consumption, at any given time, is the sum of the power consumed by all the connection requests currently provisioned in the network.

THE WPA-LR STRATEGY

The proposed routing strategy, called Weighted Power-Aware Lightpath Routing (WPA-LR), is based on a modified version of the k-shortest path algorithm [4], and it works as follows. When a connection request arrives at the network edge, up to k candidate paths are computed. The algorithm accounts for the current resource usage (i.e., wavelengths) in the network, where fiber links without available wavelengths are temporarily deleted from the network topology. When computing each candidate path, each fiber link l in the network is assigned a weight (C_l) equal to:

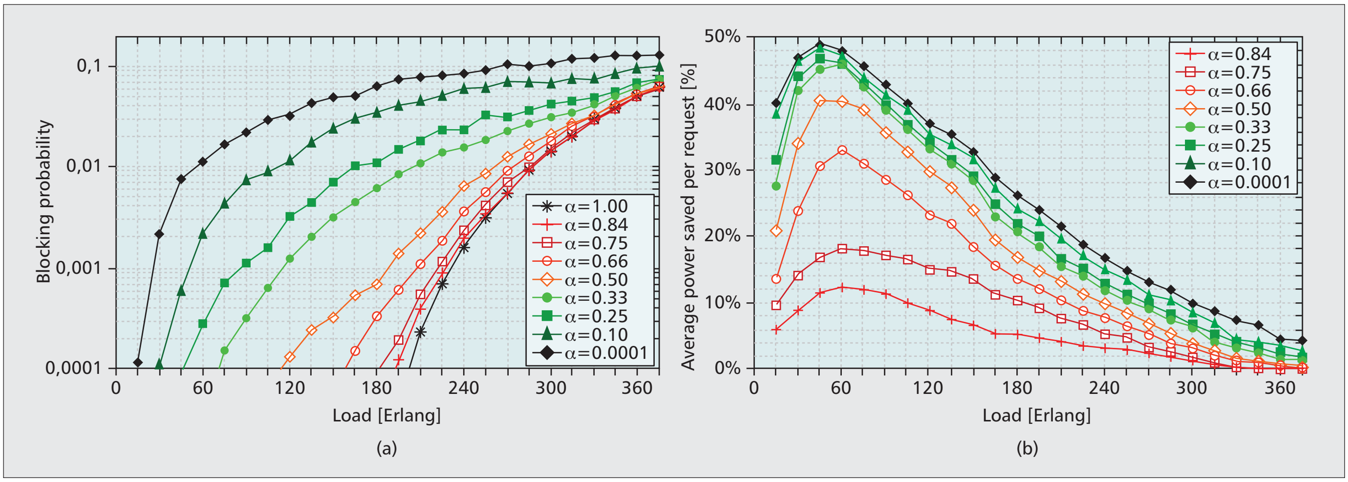

Figure 2. *a) network blocking probability; b) average power saved per request – COST 239 network.*

$$C_l = \begin{cases} \alpha \cdot P_{link,l}, \text{fiber link } l \text{ in use} \\ P_{link,l}, \text{fiber link } l \text{ not in use} \end{cases}$$

where $P_{link,l}$ represents the power necessary to operate the in-line power amplifier(s) along fiber link l, and α is a weighting factor with values between 0 and 1. Note that with values of α equal to 0 WPA-LR behaves as a pure power minimization approach, while for values of α close to 1 WPA-LR tends to provision connection requests according to shorter routes.

As a result of this first phase k paths are sorted in ascending order based on their power consumption. At this point, the first computed path satisfying the wavelength continuity constraint is chosen to provision the connection request. If, no continuous wavelength can be found among the k computed paths, or no candidate path exists, the connection request is blocked.

TRADEOFF ASSESSMENT

This section presents the impact that the WPA-LR strategy has on a series of network performance parameters, e.g., blocking probability, average fiber link utilization, path length. The results shown here are obtained via extensive simulations considering two network topologies: COST 239 [11] (with 11 nodes and 26 bidirectional fiber links), and NSFNET [12] (with 14 nodes and 21 bidirectional fiber links). It is assumed that each fiber carries a maximum of sixteen wavelengths. It is also assumed that wavelength conversion is not available in the network.

Source — destination pairs for the incoming lightpath connection requests are uniformly chosen among the networks nodes. Connection request arrivals follow a Poisson distribution while the service time for each connection is exponentially distributed. The load varies from 15 to 375 Erlangs in the experiments for the COST239 network topology, and from 15 to 210 in the experiments for the NSFNET topology. These values are chosen to investigate low, medium and high load conditions within each network. Simulation results are averaged over a series of experiments to achieve a confidence interval of 10 percent or better, with a confidence level of 90 percent.

It is assumed that EDFAs, transceivers and OXCs can be put in sleep mode immediately when not in use and that they can be turned on in a negligible transient time. It is also assumed that the process of setting to sleep mode and to waking up components does not require any additional energy, i.e., signaling issues are ignored in this work.

The following assumptions are made for the power consumption figures of the optical components: EDFAs — placed every 80 km — consume 12 W each, a transceiver 7 W (for a transmission speed of 10 Gb/s), and an OXC 6.4 W. These values are obtained, for each component, by averaging power consumption figures from commercially available data sheets. This study does not include any power consumed by equipment at higher layers different than the WDM layer (e.g., electronic layer). In this study it is also assumed that the amount of power consumed by a network element in sleep mode is negligible. This was done to provide an upper bound for the power savings that can be achieved.

For the presented results a series of values for the power-weighting factor α in the range between 10^{-4} and 1 are considered. In this study no more than $k = 3$ candidate paths are computed for each connection request, while wavelengths are assigned using a First Fit approach. The choice of $k = 3$ was based on the observation that benefit of having multiple candidate paths was obvious when increasing k from 1 to 3 while for $k > 3$ the reduction of blocking probability was negligible. On the other hand, the simulation time increased significantly with k.

POWER MINIMIZATION VS. BLOCKING PROBABILITY

The first set of results presents a trade-off between the power consumed to provision the incoming connection requests and the blocking probability. Figures 2a and 2b show the values of the network blocking probability, and the percentage of the average power saved per request,

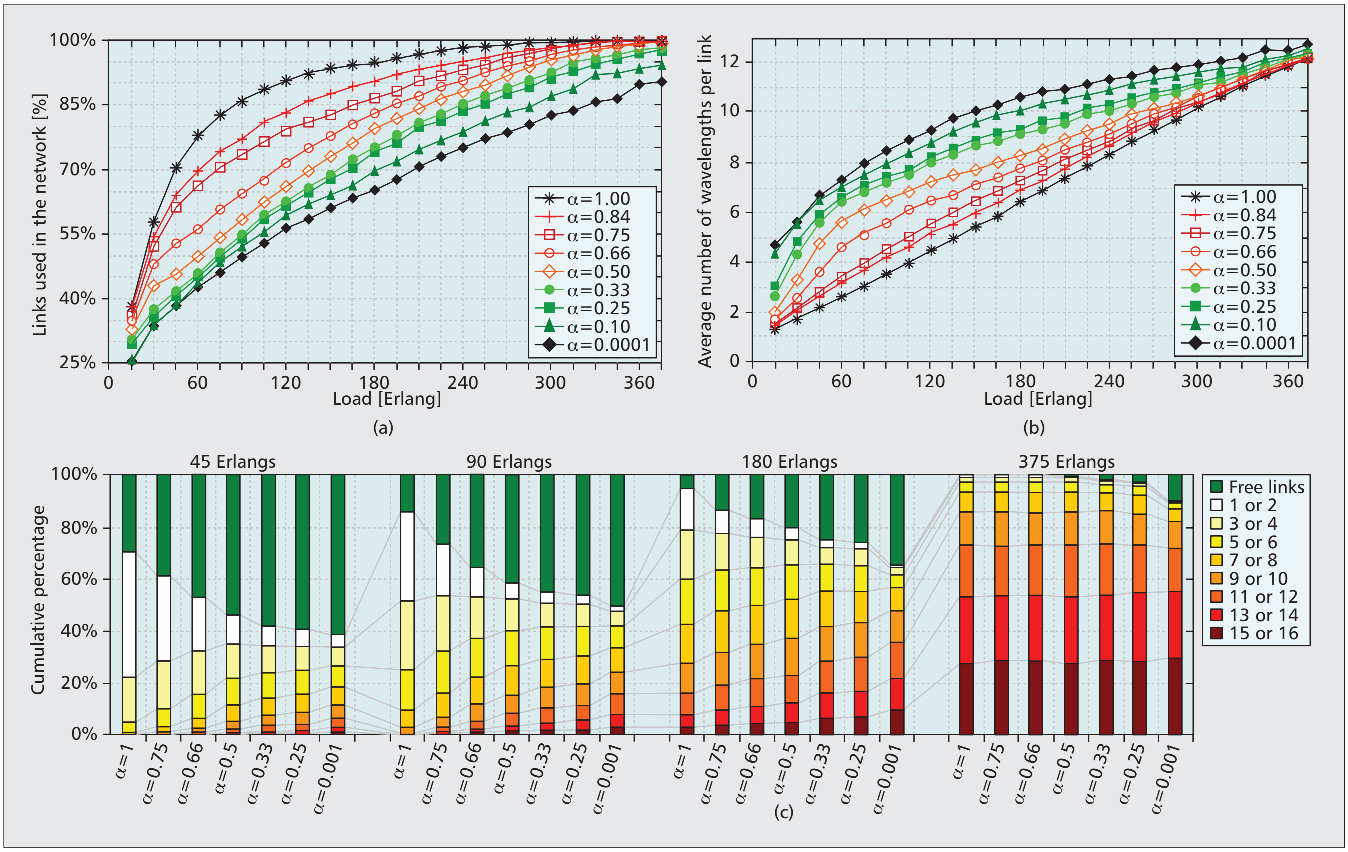

Figure 3. *Average fiber link utilization statistics – COST 239 network.*

respectively. Both metrics are presented as a function of the network load.

For any given load value, the average power per request is defined as the ratio between the total network power consumption and the number of provisioned connection requests. The average power saved per request is computed as the difference between the total network power consumption obtained when α = 1 and the total network power consumption for any other given value of α.

The figures show that considerable power savings (up to 50 percent) are achievable, but at the expense of a significant increase in the network blocking probability. This confirms that an approach that accounts only for the minimization of the network power consumption might lead to unacceptable performance degradation. This is especially true in a resource limited network scenarios, where network blocking cannot be ignored. The figures also present the possible trade-off between lower operational expenditure (OPEX) due to the reduction of the power consumption and the (possible) loss of revenue due to a higher blocking probability. For example with values of α between 0.66 and 1, there is no significant impact on the blocking probability while the power saved per request is between 30 percent and 15 percent for low and medium traffic conditions. An explanation for the observed trade-off is provided next, where the analysis of the impact of the power-weighting factor α on the fiber link utilization and on the path length values is presented.

FIBER LINK UTILIZATION

Figures 3a and 3b show the percentage of fiber links used in the network and the average number of wavelengths used per fiber link as a function of the network load, respectively. As expected the power saving figures presented previously are the consequence of an efficient utilization of network elements that are already powered-on. In Fig. 3a it can be observed that when the WPA-LR strategy focuses solely on the minimization of the length of the provisioned paths (i.e., values of α close to 1) the fiber links utilization raises very quickly with an increased value of the load. This is reasonable since, with this value of α, the routing strategy is oblivious of the power status of the network elements, and cares only about guaranteeing the shortest available path to the provisioned connection requests.

This behavior can be observed already at low traffic conditions, where more than 75 percent of the network fiber links are in use to accommodate the traffic. Conversely, when the value of α is close to 0 (i.e., when the WPA-LR strategy tries to minimize the power consumption by limiting the number of the powered on elements) the fiber link utilization increases almost linearly with the load. This translates into more fiber links (and consequently more EDFAs) that are left unused in the network.

The difference in the way fiber links are utilized has also an immediate impact on the fiber link occupancy, as shown in Fig. 3b. Fiber links are on average less utilized when α is close to 1, since with this value of α WPA-LR tends to use

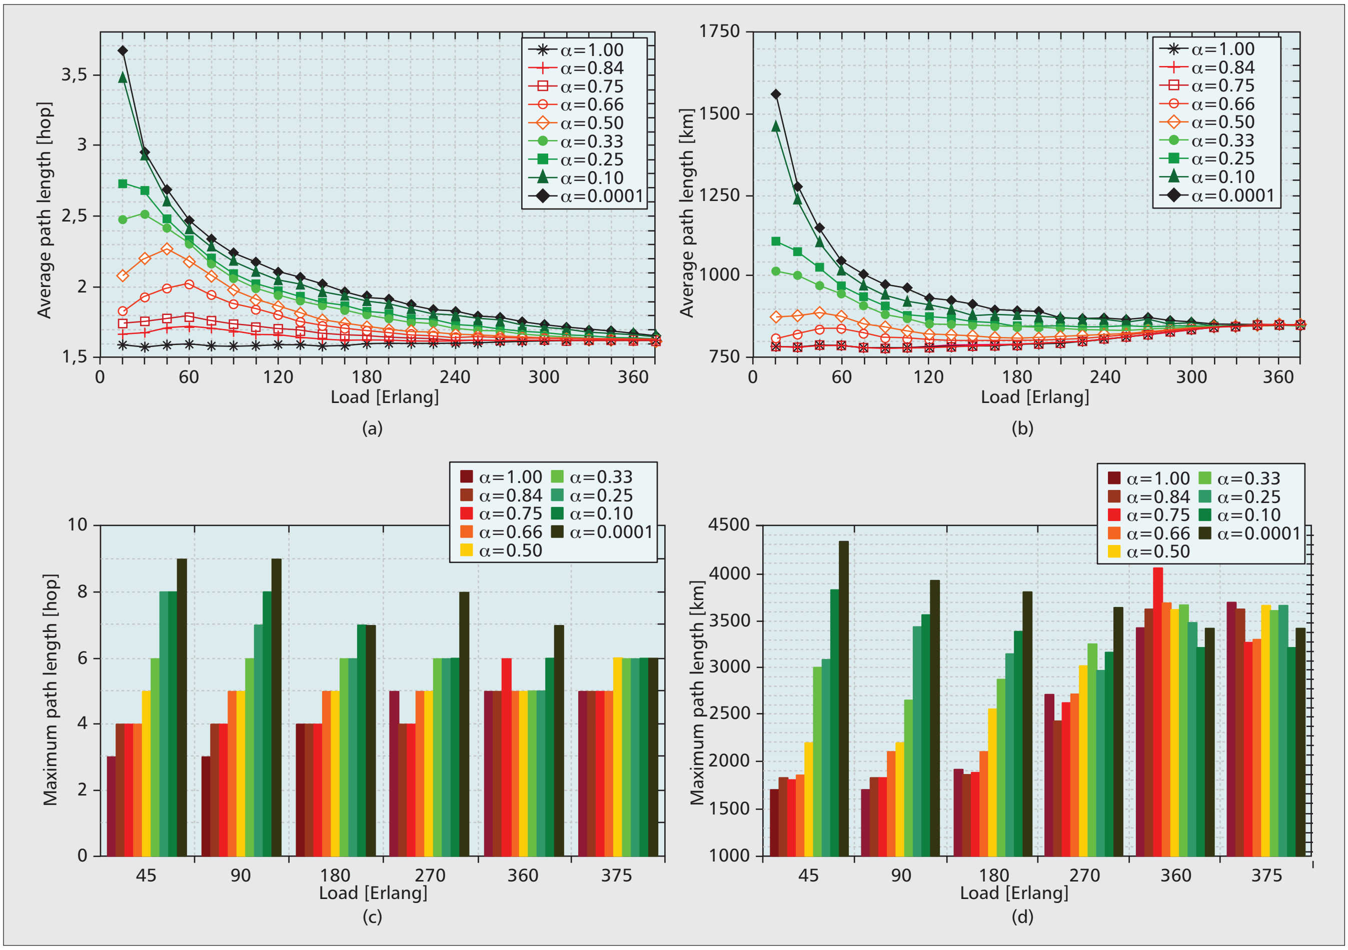

Figure 4. *Average path lengths statistics – COST 239 network.*

the shortest available path to provision a connection request. As a result the traffic is spread more evenly over the network and the chance to have a bottleneck is decreased. In contrast, when α gets closer to 0 the WPA-LR strategy tends to provision as much as possible connection requests using network elements already powered. This results in fiber links that are on average more utilized, with a higher chance to create bottlenecks, and increased blocking probability values. It can also be noted that with higher number of occupied wavelengths on some of the fiber links, the robustness of the network is reduced since it may be more difficult to reroute all the affected lightpaths upon a link failure.

Figure 3c groups fiber links based on the number of their active wavelengths. This figure confirms the consideration made above on the fiber link utilization. It shows that with any power-aware optical routing strategy (α values close to 0) the percentage of lightly loaded fiber links (e.g., the ones with up to four active wavelengths) tends to decrease, while the percentage of medium and highly loaded fiber links (e.g., the ones with 5 or more active wavelengths) increases.

PATH LENGTH

This set of results investigates the impact power-aware strategies have on the path length of the provisioned connection requests. Figures 4a and 4b present the average path length, while Figs. 4c and 4d show the maximum path length. For each set of figures both the number of hops and the geographical distance values are provided as a function of the network load. The presented results show that with a power-aware routing strategy the average and maximum path lengths are higher than the ones resulting from any power un-aware approach. This can be expected since longer paths are more likely to efficiently use network resources already powered-on. A trade-off similar to the one presented in Figs. 1a and 1b can be also observed here. For values of α between 0.66 and 1 the average path length increase is limited to 6.6 percent, an acceptable value considering the achieved energy savings.

It can also be noted that an increase in the path length degrades the signal quality at the receiver. Obviously, the impact can differ depending on length of spans, fiber type, optical amplifier and fiber characteristics, and so on. Moreover, connection reliability performance can be also affected, as longer paths might be more susceptible to experience fiber cuts.

TRADEOFF RESULTS FOR THE NSFNET TOPOLOGY

In order to be able to draw more general conclusions and confirm the validity of the findings observed so far we conduct simulation experi-

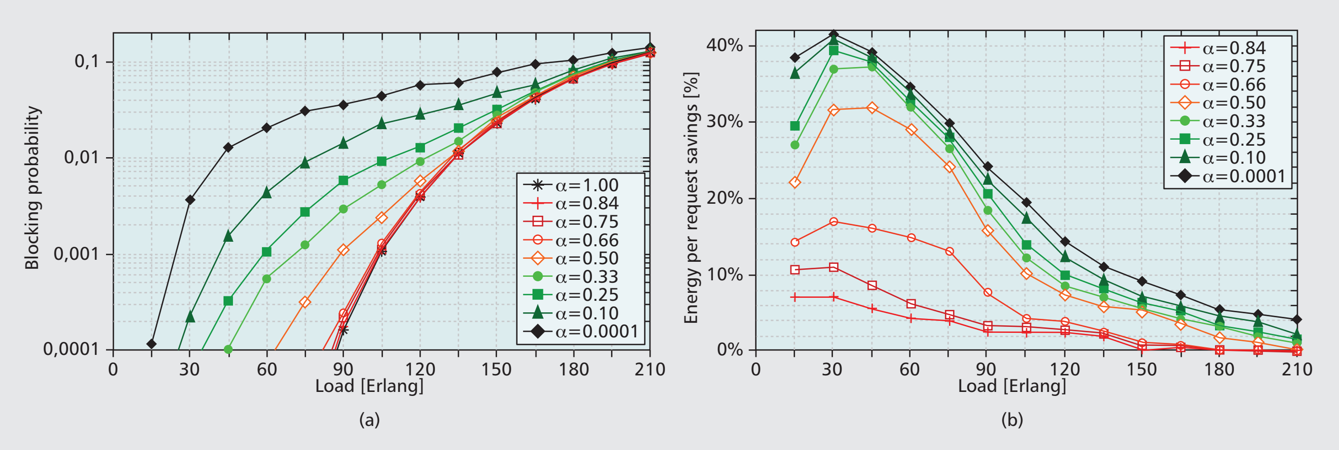

Figure 5. *a) network blocking probability; b) average power saved per request: NSFNET.*

Metrics	COST 239	NSFNET
Load with maximum power saving	45 Erlangs	30 Erlangs
Maximum power saving	49%	42%
Average path length increase (km)	46%	52%
Maximum path length increase (km)	155%	158%

Table 1. *Comparison of performance results for the COST 239 and the NSFNET network topology in load conditions where the maximum energy per request savings is observed.*

ments using the NSFNET network topology [12]. To obtain a fair comparison with the results obtained with the COST 239 network topology, the range of load values is changed to guarantee comparable performance figures in terms of blocking probability.

From the results presented in Fig. 5 it can be noticed that in general the WPA-LR strategy presents the same trade-off between power minimization and blocking probability. A similar conclusion can be drawn for the trend of the other network performance metrics summarized in Table 1.

Two main differences, on the other hand, are evident. The blocking probability seems to experience a faster increase, and maximum energy savings are observed at lower load values. One possible explanation for this behavior is the lower average nodal degree of the NSFNET topology compared to the COST239 giving a lower number of candidate paths to support energy efficient routing.

CONCLUSIONS

Although it is obvious that transparent WDM networks are promising candidates to reduce the power consumption of telecommunication networks, energy can be further saved by a proper WDM network design and a power aware lightpath provisioning. While investigating the green provisioning problem in WDM networks it was found that power minimization and some network performance metrics, e.g., blocking probability, are in conflict. The reason lays in the length of the paths originating from a pure power-aware provisioning approach. They may be too long thus increasing the network resources fragmentation level, with a detrimental effect on the network blocking probability.

The article addresses this specific issue and investigates an approach, called Weighted Power-Aware Lightpath Routing (WPA-LR), which jointly considers power minimization and resource blocking in a single cost function. Simulation results confirm the advantage of the proposed approach in terms of both power and resource usage efficiency and suggest that a solution using only power related information can lead to a significant degradation of the network performance.

It must be noted that the article does not provide a "one solution fits all" strategy. On the other hand, WPA-LR is instrumental to offer interested parties (e.g., network operators) the possibility to fine-tune the power consumption in their networks taking into account an acceptable level of blocking probability.

REFERENCES

[1] K. Sato, "Optical Technologies that Enable Green Networks," *Proc. ICTON*, 2010.
[2] A. Tzanakaki *et al.*, "Dimensioning the Future Pan-European Optical Network with Energy Efficiency Considerations," *IEEE/OSA J. Opt. Commun. and Net. (JOCN)*, vol. 3, no. 4, Apr. 2011, pp. 272–80.
[3] Y. Wu *et al.*, "Power-Aware Routing and Wavelength Assignment in Optical Networks," *Proc. ECOC*, 2009.
[4] M. Ajmal *et al.*, "Energy-Efficient Design of WDM Networks with Protection Resources in Sleep Mode," *Proc. GLOBECOM*, 2010.
[5] C. Cavdar, F. Buzluca, and L. Wosinska, "Energy-Efficient Design of Survivable WDM Networks with Shared Backup," *Proc. GLOBECOM*, 2010.
[6] I. Cerutti, N. Sambo, and P. Castoldi, "Distributed Support of Link Sleep Mode for Energy Efficient GMPLS Networks," *Proc. ECOC*, 2010.
[7] B. G. Bathula, M. Alresheedi, and J. M. H. Elmirghani, "Energy Efficient Architectures for Optical Networks," *Proc. London Commun. Symp.*, 2009.
[8] A. Jirattigalachote *et al.*, "Dynamic Provisioning Strategies for Energy Efficient WDM Networks with Dedicated Path Protection," *Optical Switching and Networking (OSN) J.*, special issue on Green Communication and Networking, vol. 8, no. 3, July 2011, pp. 201–13.
[9] Energy Star, "Small Network Equipment," http://www.energystar.gov.

[10] R. Bolla *et al.*, "Enabling Backbone Networks to Sleep," *IEEE Network Mag.*, special issue on "Energy-Efficient Networks," vol. 25, no. 2, Mar.–Apr. 2011, pp. 26–31.
[11] P. Batchelor *et al.*, "Study on the Implementation of Optical Transparent Transport Networks in the European Environment-Results of the Research Project COST 239," *Photonic Network Communications*, vol. 2, no. 1, 2000, pp. 15–32.
[12] R. Hülsermann *et al.*, "A Set of Typical Transport Network Scenarios for Network Modelling," *Photonische Netze: Vorträge der 5*, ITG-Fachtagung, vom 3. bis 4, 2004.
[13] M. Gupta and S. Singh, "Greening of the Internet," *Proc. SIGCOMM, 2003*.
[14] J. Chabarek *et al.*, "Power Awareness in Network Design and Routing," *Proc. IEEE INFOCOM*, 2008.
[15] J. C. C. Restrepo, C. G. Gruber, and C. Mas, "Energy Profile Aware Routing," *Proc. ICC*, 2009.
[16] C. Lange and A. Gladisch, "Energy Efficiency Limits of Load Adaptive Networks," *Proc. OFC/NFOEC*, 2010.
[17] F. Idzikowski *et al.*, "Saving Energy in IP-over-WDM Networks by Switching Off Line Cards in Low-Demand Scenarios," *Proc. ONDM*, 2010.
[18] Z. Yi *et al.*, "Energy Efficiency in Telecom Optical Networks," *IEEE Commun. Surveys & Tutorials*, vol. 2, no. 4, 2010, pp. 441–58.

BIOGRAPHIES

PAWEL WIATR (wiatr@kth.se) received his Master of Sciences and Engineering degree in telecommunications from National Institute of Applied Sciences in Lyon, France (INSA de Lyon) in 2001. Currently he works toward his Ph.D. degree at the Royal Institute of Technology (KTH), Stockholm, Sweden. His research interests include optical networks, energy-efficient optical networks, and optical network management. From 2005 till 2011 he worked at Comarch SA, Poland as OSS Solution Manager where he designed operation support system (OSS) solutions for network operators. He also worked as network administrator, research engineer at KTH Sweden in domains of wireless networks (IEEE 802.11b), interactive workspaces, and ubiquitous computing. Also he worked as a consultant at the University of E. Mondlane, Maputo, Mozambique.

PAOLO MONTI (pmonti@kth.se) received a Laurea degree in electrical engineering (2001) from the Politecnico di Torino, Italy, and a Ph.D. in electrical engineering (2005) from the University of Texas at Dallas (UTD). From 2006 to 2008 he worked as a research associate of the Open Networking Advance Research (OpNeAR) Laboratory at UTD. He joined KTH in September 2008 where he is currently an assistant professor in the School of Information and Communication Technology (ICT) and a member of the Optical Networks Laboratory. He has co-authored more than 40 papers published in international journals and presented at leading international conferences. He is serving on the editorial board of the Springer *Networking and Electronic Commerce Journal*. His research interests include: network planning, protocol design, performance evaluation, and optimization techniques for both optical and wireless networks.

LENA WOSINSKA (wosinska@kth.se) received her Ph.D. degree in photonics and Docent degree in optical networking from KTH, where she is currently a full professor heading the Optical Networks Laboratory and coordinating a number of national and international research projects in the School of ICT. Her research interests include fiber access networks, energy-efficient optical networks, photonics in switching, optical network management, reliability, and survivability. She has been involved in many professional activities including guest editorship of OSA, Elsevier, and Springer journals, serving on the TPCs of several conferences, as well as a reviewer for many journals and project proposals. She was a Co-Chair of SPIE APOC 2008 and IEEE/OSA/SPIE ACP 2009 Subcommittee on Network Architectures, Management and Applications (SC4), and Chair of SC4 in IEEE/OSA/SPIE ACP 2010 and 2011. She was General Chair of the IEEE International Conference on Transparent Optical Networks (ICTON 2011) and is currently Chair of the Optical Networks and Systems Symposium within IEEE ICC 2012. In 2007–2009 she was an Associate Editor of *OSA Journal of Optical Networking*, and since April 2009 she is serving on the Editorial Board of *IEEE/OSA Journal of Optical Communications and Networking*.

> *Although it is obvious that transparent WDM networks are promising candidates to reduce the power consumption of telecommunication networks, energy can be further saved by a proper WDM network design and a power aware lightpath provisioning.*