On the Energy Efficiency of IP-over-WDM Networks

Francesca Vismara, Vida Grkovic, Francesco Musumeci, Massimo Tornatore, Stefano Bregni
Politecnico di Milano, Department of Electronics and Information, 20131 Milan, Italy
E-mail: fmusumeci@elet.polimi.it

Abstract*—In the next 10 to 15 years the Internet will undergo a substantial increase especially with respect to the bandwidth required by end-users. Since the current Internet already consumes a not-negligible percentage of the total world electricity, reducing the energy consumption of telecom networks is expected to become an increasingly-important challenge, being unacceptable that the Internet energy consumption grows proportionally to the served bandwidth. In this paper we focus on backbone transport networks, that serve large aggregated amount of traffic. We compare three different network architectures which implement the Wavelength Division Multiplexing (WDM)-based transport of IP packets over optical fiber links (IP-over-WDM networks). The differences between these three architectures, which we identify as *IP with no Bypass* (*IP-NB*), *IP with Bypass* (*IP-B*) and *IP with Bypass and Grooming* (*IP-BG*), concern the capabilities of performing aggregation of traffic (grooming) and optical switching. IP-NB architecture performs grooming in every network node, where the traffic is electronically processed and forwarded by the IP routers. IP-B enables switching of wavelength channels directly in the optical domain, thus bypassing the processing of IP packets in the intermediate IP routers. This architecture does not provide grooming capabilities, but it just allows aggregation of different traffic demands to be established between the same source/destination pairs. IP-BG architecture represents an intermediate solution between the previous, since it provides both grooming capabilities (as for IP-NB) in order to efficiently exploit network capacity, and optical switching (as for IP-B) to reduce expensive electronic processing operations. We perform a comparative study between these three architectures showing the trade-off between the reduction of the power consumption or the cost of the networks, and we analyze how minimizing one of these two factors can influence the other.

I. Introduction

The bandwidth growth that the Internet will face in the near future represents a very challenging issue from the power consumption point of view. Nowadays, the 7-8% of the world energy consumption is absorbed by the ICT [1] and specifically the Internet is responsible for about the 25% of this amount. Moreover, it is estimated that during the next 10 to 15 years, the traffic-bandwidth requirement will be up to 50-times higher than the current one. This is mainly due to the evolution of the Internet towards a home/mobile user network service where most of the traffic consists of high-speed video-based streams. In this scenario, it is quite intuitive that the power requirement will be the major constraint for the next generation networks.

We focus on optical networks for backbone transport of traffic, and we consider the reduction of their power consumption as well as of their cost, enabled by novel optical networking technologies. Specifically, optical networks employing the Wavelength Division Multiplexing (WDM) technique are capable to route wavelength channels (i.e., *lightpaths*), each of them carrying one (or more) optical signal(s), from source to destination nodes, through a series of optical fiber links and, in most of cases, with no need to perform data processing, i.e., reducing the Optical/Electronic/Optical (O/E/O) conversions of the signal (which require high power consumption due to the opto-electronic devices and electronic processing needed). In IP-over-WDM (IPoWDM) networks, a virtual (logical) topology of optical circuits carrying IP-based traffic is established upon a physical topology made up of WDM links, and these optical circuits are maintained in the optical domain, e.g., from a source IP router port up to destination IP router port. IPoWDM architectures can have different implementations: a possible approach consists in terminating optical circuits at every node, thus using optical technology just for the transmission of the signal over WDM fiber links (*IP with no Bypass* [2], or *IP-NB*). This technique allows aggregation of traffic from multiple sources into the same optical channel (*grooming*) and separation of the optical channel traffic towards different destinations (*degrooming*), but implies expensive and energy-requiring operations such as O/E and E/O conversions and electronic processing. An alternative solution consists in interconnecting all the IP routers in the network by means of direct optical circuits (*IP with Bypass* [2], or *IP-B* architecture). However, this approach is not efficient from the network capacity point of view since each connection request, even if it requires a small bandwidth, needs a *whole* dedicated λ (i.e., an entire wavelength channel). *IP with Bypass and Grooming* (*IP-BG*) architecture represents a trade-off between the previous two solutions: each intermediate node traversed by a connection request (i.e., a node different from the source/destination pair of that request) can be directly bypassed by the lightpath, which can be switched directly in the optical domain, or can provide grooming capabilities by performing electronic processing at IP level, exploiting the bandwidth capacity more efficiently.

Since most of today's traffic demands requires lower bit-rates than those provided by a full wavelength channel, it is natural to think about how to effectively pack traffic demands in the large bandwidth of wavelength channels. The most commonly adopted strategy to solve the problem of carrying several subrate traffic demands in a single wavelength is traffic grooming: it allows demands to be transported together on the

978-1-4244-7173-7/10/$26.00 ©2010 IEEE

same wavelength even though their source/destination node pairs are not the same.

Various very recent works have investigated on the effect of traffic grooming over the energy efficiency of the transport layer of optical networks. In [2] the authors exploit the concept of lightpath-bypass in order to design an IPoWDM network that minimizes power consumption by reducing the number of needed IP router ports. The approach developed in [3] models the power consumption of an individual lightpath; then, the overall power consumption of the whole network can be obtained as a function of the number of established lightpaths, which strictly depends on how traffic grooming is performed. In [4], two approaches are considered to formulate the power consumption of the network according to the fact that: i) power consumption is roughly considered as a linear function of the traffic load, or ii) that the impact of traffic load over the network power consumption can be neglected. Modular nodes, consisting in two main sections, i.e., photonic and electronic, have been considered in [5] to perform energy-aware dynamic traffic grooming in order to gain substantial energy-savings in network operations. Traffic grooming is also considered in [6] as a means to reduce operational overhead and to let the power consumption be proportional to the bandwidth requirement.

In this paper, for the first time to the best of our knowledge, we consider three different IPoWDM architectures and we explore the trade-offs in terms of power consumption and cost among the three aforementioned architectures. In the results section we numerically illustrate how minimizing one of these two factors can influence the other.

The rest of the paper is organized as follows: in section II we discuss all network components that contribute to the overall network power consumption, and we classify the corresponding contributions. Similarly we show the normalized costs related to the considered network components. In section III the three transport network architectures are described with more detail. Section IV presents simulation results which minimize the overall network power consumption of the three architectures as well as results obtained by minimizing the network cost. Finally, we conclude the paper in section V, where we summarize our results.

II. A Classification of Power and Cost Contributions in IP-over-WDM Networks

A. Power Contributions

Several components can influence the total power consumed at the transport level of an IPoWDM network, e.g.:

- switching devices (either electronic or optical);
- devices for transmission/reception of optical signal, that we generically refer to as transponders (intended as IP router ports, i.e., transmitter or receiver cards);
- pumping of optical amplifiers (usually Erbium Doped Fiber Amplifiers or EDFAs);
- network signaling.

However, in a first approximation, some of these contributions can be considered constant through the three architectures under analysis, i.e., they are the same in all the models: e.g., the power contribution related to the signaling operations performed by the control plane can be assumed to be constant (equal to about 150 W for a generic transport network [7]). Moreover, since offered traffic must be processed in its source and destination IP routers, independently of the transport architecture being considered, this contribution also requires a fixed power consumption and only the packet processing performed in the intermediate nodes will be taken into account. Finally, each EDFA needs a pump power to perform the amplification of all the optical signals carried by the different wavelengths (it requires about 4.5 W per amplifier per fiber according to [2]). We opted here for a model in which the EDFAs are given and already placed every 80 km [2], so that the total contribution due to EDFAs can be also considered a constant value for all the network architectures. In conclusion, the remaining variable contributions to total power consumption can be grouped into the following three categories:

1) *Transponders*: each source router port is equipped with a transponder in order to convert the electronic signal into an optical signal and to transmit it over one of the available wavelengths; on the other hand each destination router port needs a transponder to receive the signal and convert it back into the electronic domain. Moreover, a couple of transponders (receiver and transmitter) is also necessary in intermediate nodes in case traffic grooming/degrooming is performed.
2) *Electronic processing*: this relevant contribution to power consumption arises when electronic processing is performed in intermediate nodes, typically to accomplish grooming and degrooming of traffic demands concerning different source/destination pairs. It is estimated [8] that the power consumption of an IP router (used, for example, to perform grooming) is on the order of 17.5 W per Gbit/s of processed traffic.
3) *Optical switching*: switching can be also performed directly in the optical domain by the optical switching fabrics. Such devices are usually based over MEMS (Micro-Electro-Mechanical-Systems) mirrors and this technology is here considered. When a mirror of a 2D MEMS-based optical switch is actuated (that is, whenever a wavelength channel is switched), a power on the order of 11.5 mW is required [9].

In the following the overall contributions due to transponders, electronic processing and optical switching will be referred to as P_{tr}, P_e and P_o, respectively.

B. Cost Contributions

Similarly to the previous power consumption analysis, we consider the same three contributions (each of them corresponding to a power contribution) also for the cost comparison. All costs are normalized to a unitary value, thus we compare the three architectures in terms of total normalized cost. In [10] a set of realistic cost values are presented, without any reference to specific vendors products. Specifically, the three cost components are the following:

1) *transponders*: 10 Gbit/s Long-Haul transponders are here considered and they are supposed to have a normalized cost equal to 2;
2) *electronic processing*: IP router ports, including electronic switching as well as packet/header-processing functionalities performed in intermediate nodes, is considered to have a normalized cost equal to 0.2 per Gbit/s of traffic being treated;
3) *optical switching*: a 10 Gbit/s single lightpath switching operation, performed by a MEMS-based device, is accomplished by spending a normalized cost equal to 0.05.

In Table I the considered power and normalized cost contributions are summarized.

	Power Requirement	Normalized Cost
10 Gbit/s Transponder	35 W	2
Electronic processing	17.5 W per Gbit/s	0.2 per Gbit/s
Optical switching	11.5 mW per lightpath	0.05 per lightpath

TABLE I
POWER CONSUMPTION CONTRIBUTIONS AND NORMALIZED COSTS.

III. NETWORK MODELS DESCRIPTION

In the following three IP-over-WDM architectures are compared. The differences between these models mainly concentrate on how grooming and switching operations are performed. Therefore the power requirements and the corresponding costs to be accounted for vary according to the considered technology. Our results have been obtained by using Integer Linear Programming (ILP) formulations[1] which aims at minimizing the overall network power consumption in the first case, and the network cost in the second case. The features of the three models can be summarized as follows:

- *IP-NB*: here switching and grooming are both accomplished in the electronic domain: in every network node, as shown in Fig. 1, the optical signal is converted into the electronic domain through a transponder. Then it is processed by the IP router, which forwards it towards the next network node and, if necessary, provides for grooming it with different optical signals into the same wavelength. Finally, it is converted back into the optical domain through a transponder. Thus in this scenario, there is no optical switching power contribution. The ILP model used for this architecture consists of a multi-commodity flow-based formulation [11], where a single layer topology is modelled for the routing assignment problem. This formulation is only subject to the capacity constraint, i.e., each link connecting two different nodes can transport a limited amount of traffic according to a fixed capacity, and to the flow conservation constraint, i.e., all the traffic entering each network node must be forwarded by the node (if it represents an intermediate node for the specific request) or terminated in it (in case it is the destination node of the request).
- *IP-B*: as represented in Fig. 2, at each node a router (IP layer) is placed over a MEMS-based switch (optical transport layer). In this case switching is performed in the optical domain, while grooming can be made only for the same source/destination node pairs: so the number of transponders needed strictly depends on the number of requests and on their required bandwidth (i.e., the total number of lightpaths needed). The ILP here considered is a single layer-based Routing and Wavelength Assignment (RWA) problem [12], in which we need to consider the so-called λ-continuity constraint in addition to those already described for the IP-NB architecture. Since each connection request uses one (or more) whole lightpath(s) we need to assign a (set of) wavelength(s) to it. Moreover, we assume that optical switches do not have wavelength conversion capabilities, thus each lightpath must be associated to the *same* wavelength in all network links it traverses.
- *IP-BG*: this network architecture is an intermediate one between IP-B and IP-NB. Grooming and therefore O/E/O conversion as well as electronic processing, are performed only where necessary. The power contributions considered here, as shown in Fig. 3, are due to electronic processing performed in the intermediate nodes (where grooming is accomplished, as in IP-NB case), optical

[1]The ILP formulations are here omitted, but we will present them in a future work.

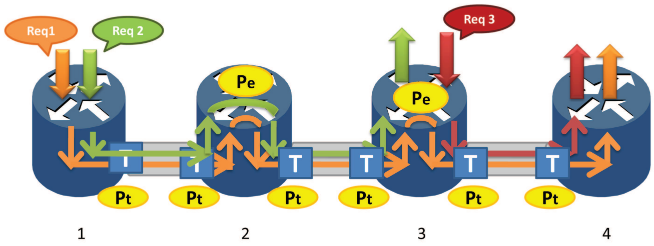

Fig. 1. *IP-NB* network architecture: two requests (Req1 and Req2) are groomed and transported by the same lightpath between nodes 1 and 2, where the whole traffic is processed by the IP router, and between nodes 2 and 3, which is the destination node of Req2. Node 3 performs grooming of Req1 and Req3 which are both sent to the destination node 4.

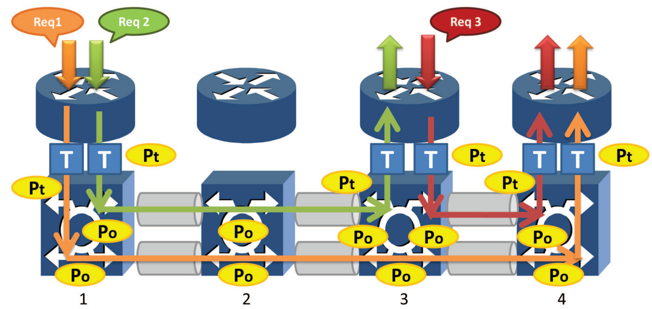

Fig. 2. *IP-B* network architecture: each request uses a whole λ (two λs are required) and its traffic is processed at the IP layer only in source and destination nodes, whereas in intermediate nodes lightpaths are optically switched by MEMS-based switching fabrics.

	IP-NB	*IP-B*	*IP-BG*
ILP model	Single layer Multicommodity-Flow-based Formulation	Single layer RWA Formulation based on [12]	Double Layer Flow Formulation based on [13]
Switching	Electronic (E)	Optical (O)	Mixed (E & O)
Grooming	Performed	Not performed	Performed
Objective	$minimize\ P_{tot} = P_{tr} + P_e$	$minimize\ P_{tot} = P_{tr} + P_o$	$minimize\ P_{tot} = P_{tr} + P_e + P_o$
Constraints	Flow Conservation and Capacity constraint	Flow Conservation, Capacity constraint and λ-continuity constraint	Flow Conservation, Capacity constraint and λ-continuity constraint

TABLE II
ILP MODELS FOR THE THREE ARCHITECTURES.

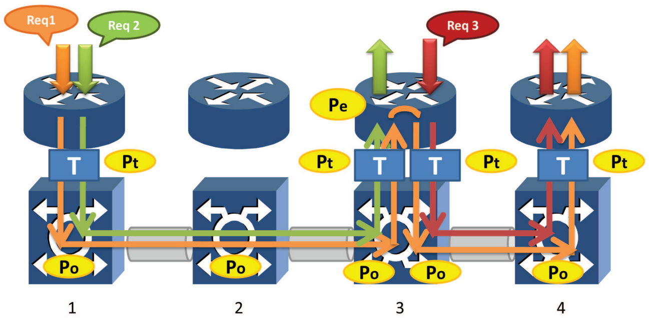

Fig. 3. *IP-BG* network architecture: the requests can be groomed together (only one λ is needed), but the traffic is electronically processed only when necessary: the IP router at node 2 is optically bypassed and at node 3 Req1 is processed in order to be degroomed by Req2 and groomed with Req3.

switching (in intermediate nodes where optical bypass is implemented, as in IP-B case), and transponders. The ILP considered for this architecture consists of a double-layer (logical+physical) flow formulation, where a logical topology (the set of the established lightpath) is mapped over a physical topology (the set of physical links). Each lightpath is routed through a set of physical links and is subject to the λ-continuity constraint already described. Moreover the flow conservation and the capacity constraints must hold.

Figures 1, 2 and 3 depict the power contributions which have to be taken into account in the three cases by routing three connection requests characterized by subrate traffic bandwidth, whereas the main features of the ILP models we have used for the three architectures have been summarized in Tab. II.

IV. SIMULATION RESULTS

In this section, we show some numerical results that assess the overall power consumption and cost for the three aforementioned transport architectures. Two different network topologies have been considered. In the first case, a traffic matrix having a total amount of traffic of 350 Gbit/s has been mapped over an EON core network topology consisting of 11 nodes and 26 bidirectional links (COST239 network in [14]). In the second case, a nonuniform traffic matrix characterized by an overall amount of traffic equal to 180 Gbit/s has been applied to be mapped over a NSFNET network topology consisting of 14 nodes and 22 bidirectional links [2].

A. Power consumption minimization

Fig. 4 shows simulation results optimized with respect to power consumption, obtained by the ILP models mentioned in Tab. II. Specifically, Fig. 4(a) shows the overall power consumed by the three network architectures, whereas Fig. 4(b), (c) and (d) show the different power contributions (P_{tr}, P_e and P_o, respectively). We can immediately observe that both in the EON and in the NSFNET case the IP-BG architecture results in the lowest power consumption, typically at least 30% lower than the IP-NB case and more than 25% lower than the IP-B case. As we can see from Fig. 4(b)-(c), this is mainly due to the reduction achieved by IP-BG on the number of transponders (as compared to both IP-NB and IP-B) and electronic processing (as compared to IP-NB, since IP-B does not have P_e).

It is worth noting that in the NSFNET case the IP-NB architecture consumes the highest power (mainly due to high power for electronic processing), while, in the EON case, the IP-B architecture is the least energy-efficient (mainly due to a very large number of transponders). This different behaviour derives from a different effect of the network topology and the overall offered traffic considered in the two cases. In fact, the EON topology is considerably more meshed than the NSFNET and it is also characterized by much lower average hop length for a connection[2]: it follows that the amount of electronic processing in the intermediate nodes for the IP-NB case is relatively much lower in the EON than in the NSFNET. Moreover, the high amount of traffic considered for the EON case (350 Gbit/s of traffic, versus 180 Gbit/s for the NSFNET case) also tends to make the IP-B architecture the least efficient among the three, since the number of transponders needed in

[2]The average number of links per node is 4.7 for the EON and 3 for the NSFNET, whereas the average shortest-path hop length is 1.56 for the former and 2.12 for the latter

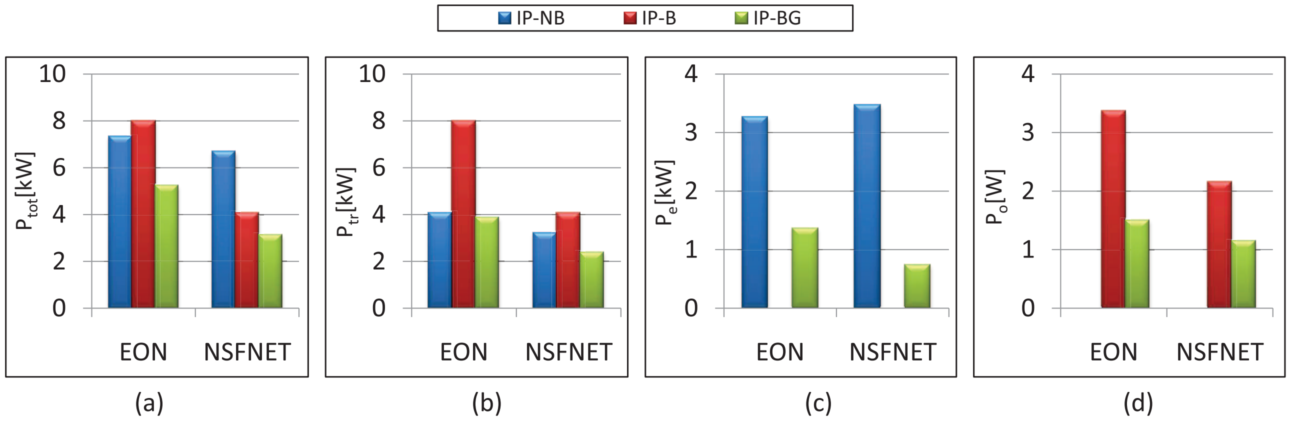

Fig. 4. (a) Total power consumption and single contributions for the three architectures: (b) P_{tr}, (c) P_e and (d) P_o.

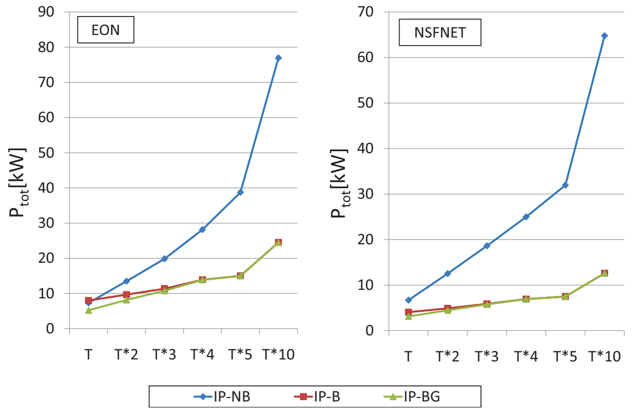

Fig. 5. Total power consumption with respect to traffic matrix scaling factor.

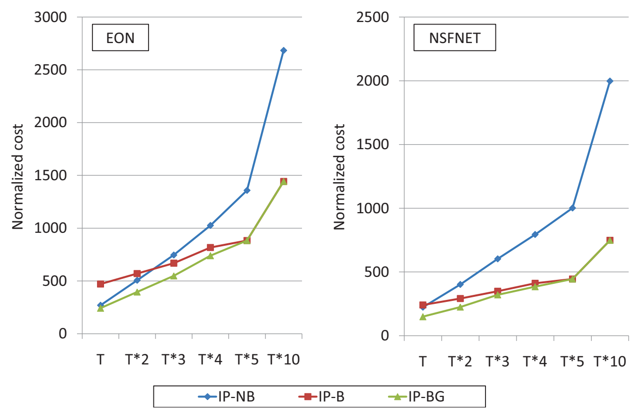

Fig. 6. Normalized costs obtained after power consumption minimization.

IP-B (which gives the main contribution to the total power consumption of this architecture) is strictly related to the bandwidth of traffic requests (while it is practically *independent* of the network topology). Therefore, given a certain traffic matrix, changing the topology of the network (e.g., adding some links between its nodes) can provide benefits only for the IP-NB (and obviously also for the IP-BG) case. Finally, the optical switching power contribution is, as we expected and as we can observe by the charts of Fig. 4(d), the lowest among the three contributions. Note that P_o is of the order of Watts, whereas P_{tr} and P_e are of the order of kW.

For both the topologies, the same analysis was carried on by considering a traffic matrix where the bandwidth of each traffic demand has been scaled by a factor of 2, 3, 4, 5 and 10, respectively. In Fig. 5 we show how the power consumption grows for increasing traffic demand in two network topologies. First, we can observe that the power requirement of IP-NB increases much more rapidly than the other two architectures. Second, for increasing traffic demands, the IP-B architecture succeeds in exploiting more efficiently the capacity of each wavelength channel, since each demand requires a bandwidth that tends to occupy the *entire* capacity of each lightpath. As a consequence, at high traffic loads (i.e., when the traffic matrix is scaled by a factor greater than or equal to 4 in the EON case and 3 for the NSFNET case), the IP-BG architecture tends to behave as the IP-B one, avoiding to perform grooming and assigning to each traffic demand the capacity of one (or more) entire wavelength(s).

We now derive the total costs of the three architectures when the EON and NSFNET topologies are designed in order to minimize power consumption: results show them in Fig. 6, and are obtained by using the normalized costs described in section II, The costs show that, when we aim at minimizing network power consumption, the IP-BG solution is not only the most power-efficient architecture among the three, but also the most cost-effective one. At low traffic loads (especially for the EON case) IP-NB is almost as cheap as IP-BG, but as traffic increases, IP-NB becomes much more costly than IP-BG and IP-B.

B. Network cost minimization

Another problem was also considered in this study: we have designed the network in order to minimize cost, and then observed the corresponding power consumption. In Fig. 7 results are shown for the two network topologies for different traffic-load scaling factor. As for the power minimization problem, we

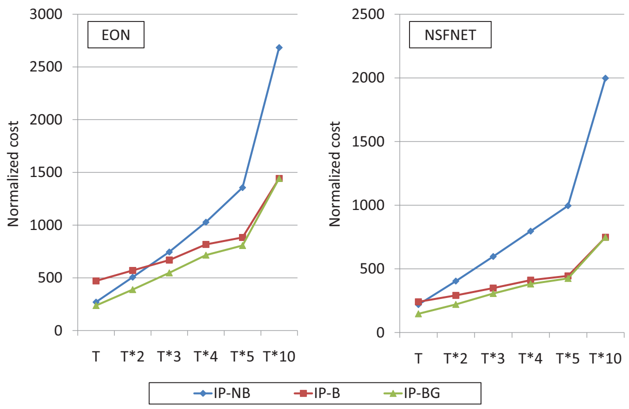

Fig. 7. Total network cost for increasing traffic-matrix scaling factor.

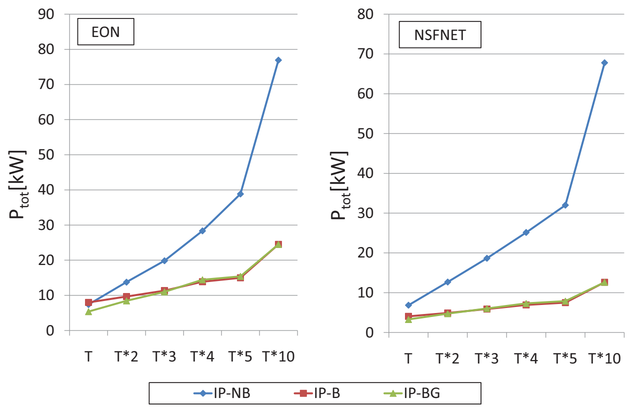

Fig. 8. Power consumptions obtained after cost minimization.

observe that IP-BG is the most efficient architecture and IP-NB is the most expensive one. At low traffic loads, IP-B is a costly solution due to the high number of transponders needed, but as the traffic bandwidth scaling factor increases, this architecture exploits wavelengths capacity in a more efficient manner, becoming a cost-effective solution.

After the total network cost is minimized for the three architectures, we observe the corresponding power consumptions and show them in Fig. 8. Even when aiming at minimizing network cost, IP-BG is the most cost-effective architecture, whereas IP-NB is the most expensive one, especially at high traffic loads, when the difference between the cost of IP-NB and the other two architectures becomes even higher.

The obtained results have shown that the design of an energy-efficient network is in agreement with a cost-effective network design.

V. Conclusion

In this paper we have compared three different transport network architectures with respect to power consumption and cost. Grooming capabilities provided by network nodes improve power savings by efficiently exploiting network capacity, and thus, network resources. We have used ILP formulations to compare three different transport network architectures in order to minimize the power consumption or the cost of the network. Simulation results have shown that the IP-BG architecture is the most efficient solution from both the power consumption and the cost point of view. At low traffic loads the IP-B architecture has the worst performance among the three since it does not efficiently exploit network resources, whereas at high traffic loads it provides power and cost savings if compared to the IP-NB architecture. IP-BG is an intermediate solution between the others, and in fact its behaviour at low traffic loads is similar to the IP-NB one (i.e., it performs grooming more frequently), but at high traffic loads it is similar to the IP-B one (i.e., it does not perform grooming since each request is capable to occupy the *whole* capacity of the required wavelengths). Finally, it is also interesting and useful to notice that an energy-efficient network design is also a cost-efficient design, especially because IP router ports play a dominating role in IP-over-WDM networks from both the energy consumption and network cost point of view.

References

[1] W. Vereecken, L. Deboosere, D. Colle, B. Vermeulen, M. Pickavet, B. Dhoedt, and P. Demeester, "Energy Efficiency in Telecommunication Networks," in *European Conference on Networks and Optical Communications & Optical Cabling and Infrastructure (NOC08)*, Krems, Austria, Jul. 2008.

[2] G. Shen and R. Tucker, "Energy-Minimized Design for IP Over WDM Networks," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 1, no. 1, pp. 176–186, Jun. 2009.

[3] E. Yetginer and G. Rouskas, "Power Efficient Traffic Grooming in Optical WDM Networks," in *Global Telecommunications Conference, IEEE GLOBECOM 2009*, Honolulu, Hawaii, Nov. 2009, pp. 1–6.

[4] S. Huang, D. Seshadri, and R. Dutta, "Traffic grooming: a changing role in green optical networks," in *Global Telecommunications Conference, IEEE GLOBECOM 2009*, Honolulu, Hawaii, Nov. 2009, pp. 1–6.

[5] M. M. Hasan, F. Farahmand, and J. P. Jue, "Energy-Awareness in Dynamic Traffic Grooming," in *Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference, 2010 Conference on (OFC/NFOEC)*, San Diego, CA, USA, 21-25 Mar. 2010, pp. 1–3.

[6] M. Xia, M. Tornatore, Y. Zhang, P. Chowdhury, C. Martel, and B. Mukherjee, "Greening the Optical Backbone Networks: A Traffic Engineering Approach," in *Communication (ICC), 2010 International Conference on*, Cape Town, South Africa, 23-27 May 2010.

[7] S. Aleksic, "Power Consumption Issues in Future High-Performance Switches and Routers," in *Transparent Optical Networks, 2008. ICTON 2008. 10th Anniversary International Conference on*, vol. 3, Athens, Greece, 22-26 Jun. 2008, pp. 194–198.

[8] P. Chowdhury, M. Tornatore, and B. Mukherjee, "On the Energy Efficiency of Mixed-Line-Rate Networks," in *Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference, 2010 Conference on (OFC/NFOEC)*, San Diego, CA, USA, 21-25 Mar. 2010, pp. 1–3.

[9] Y. Gokdel, B. Sarioglu, and A. Yalcinkaya, "LED Integrated Miniaturized Polymer MEMS Display," in *Research in Microelectronics and Electronics, 2008. PRIME 2008. Ph.D.*, Jun. 2008, pp. 93–96.

[10] R. Huelsermann, M. Gunkel, C. Meusburger, and D. A. Schupke, "Cost Modeling and Evaluation of Capital Expenditures in Optical Multilayer Networks," *Journal of Optical Networking*, vol. 7, no. 9, pp. 814–833, 2008.

[11] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[12] B. Mukherjee, *Optical Communication Networks*. McGraw-Hill, 1997.

[13] K. Zhu and B. Mukherjee, "Traffic Grooming in an Optical WDM Mesh Network," *Selected Areas in Communications, IEEE Journal on*, vol. 20, no. 1, pp. 122–133, Jan. 2002.

[14] M. Liu, M. Tornatore, and B. Mukherjee, "New and Improved Strategies for Optical Protection in Mixed-Line-Rate WDM Networks," in *Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference, 2010 Conference on (OFC/NFOEC)*, San Diego, CA, USA, 21-25 Mar. 2010, pp. 1–3.