POLITECNICO DI TORINO
Repository ISTITUZIONALE

Traffic grooming and energy-efficiency in flexible-grid networks

(Article begins on next page)

23 April 2024

Traffic Grooming and Energy-efficiency in Flexible-Grid Networks

Arsalan Ahmad, Andrea Bianco, Edoardo Bonetto

Department of Electronics and Telecommunications, Politecnico di Torino, Torino, Italy

Email: {firstname.lastname}@polito.it

Abstract—**Energy-efficient design of flexible-grid networks is investigated. We focus on the design of the logical layer, usually disregarded when dealing with flexible-grid networks. More precisely, we evaluate the impact of introducing an energy-aware electronic traffic grooming in flexible-grid networks design. We propose two greedy heuristics for the network design, one exploiting traffic grooming, and we compare their energy efficiency. Results have been retrieved for several randomly generated networks of different size, with different connectivity, average physical link length and traffic scenarios. Significant energy savings can be achieved for low traffic loads and large network size when performing traffic grooming.**

I. INTRODUCTION

The Information and Communication Technology (ICT) sector is estimated to be responsible for a percentage between 2% and 10% of the worldwide energy consumption. Current trends predict also that the Internet will consume 50% of the world electricity in the next future [1]. Indeed, since the network traffic is expected to steadily increase in the next years, the telecommunication networks, and in particular the backbone networks, will experience a large increase of their energy consumption, which will become a critical issue from both a technical and economical point of view. At present, backbone networks consume about 20% of the total energy of the Internet and it is estimated that their power consumption could become a dominant part of the overall Internet energy requirements in a near future [2].

Thus, improving network energy efficiency is a very challenging topic, extensively addressed nowadays by researchers and manufacturers. Indeed, energy efficient devices are starting to be introduced in the market, but to achieve significantly higher energy efficiency, the most viable strategy is to introduce energy awareness during the network planning and design phases.

It has been demonstrated that introducing energy awareness in the design of Wavelength Division Multiplexing (WDM) backbone networks can improve the energy efficiency [3]. However, WDM core networks may not be the optimal solution for future backbone networks. The recently developed optical Orthogonal Frequency Division Multiplexing (OFDM) transmission has introduced high bit rate channels and a higher spectral efficiency thanks to a flexible use of the optical spectrum. These characteristics permit to cope with the increasing traffic load, while allowing a reduced energy consumption thanks to a more effective use of the deployed network resources.

The concept of flexible-grid networks has been recently widely investigated. In this paper we focus on the impact of energy awareness in the static network design of flexible-grid optical backbone networks. Indeed, flexible-grid networks are usually studied emphasizing the spectral efficiency, while only few works consider them as a solution to the energy consumption issue. Our contribution is to evaluate the impact on the energy consumption of performing an accurate energy-aware static network planning of the logical layer of the network with respect to a not power-optimized one. In particular, we want to check if performing electronic traffic grooming in flexible-grid networks can improve the energy efficiency as it has been shown for WDM networks [4]. Indeed, as discussed in Sec. II the optimization of the logical layer has been not examined in depth in previous works.

Thus, we propose two greedy heuristics which perform the complete network design, the first one mostly establishing direct lightpaths, the second one trying to optimize energy efficiency by exploiting electronic grooming when convenient. Furthermore, we perform an accurate sensitivity analysis of network's parameters, such as the link length, the node degree connectivity and the traffic load, to understand how each parameter can influence the network design.

The remaining part of our paper is organized as follows. In Sec. II, the characteristics of flexible-grid networks and the network model are introduced and a comparison with previous works is also performed. The proposed heuristics are detailed in Sec. III. In Sec. IV the network scenarios and the retrieved results are presented. Lastly, conclusions and possible directions for the future work are summarized in Sec. V.

II. FLEXIBLE-GRID NETWORKS

Current backbone networks are based on WDM systems and they follow the ITU fixed-grid frequencies for the allocation of wavelength channels. However, the fixed size frequency allocation in WDM networks has several drawbacks due to its coarse granularity and non flexible nature. In particular, the main issue is the mismatch between the provided bandwidth and the required one. Each wavelength channel provides a fixed bandwidth capacity which can not be adapted to the one required by the traffic demand. Moreover, the WDM networks will become more spectrally inefficient with the evolution of wavelength channel capacity, e.g., from current 10 and 40 Gb/s to 100 Gb/s and beyond.

A solution can be the adoption of optical OFDM transmission and the evolution of backbone networks from fixed-grid WDM to flexible-grid networks. Indeed, in flexible-grid networks the optical spectrum can be allocated in a grid-less way which means that any portion of the spectrum can be

used without being constrained by the ITU grid spacing. Furthermore, OFDM allows to adapt the transmission rates to the bandwidth requirements of the traffic demands by modifying the modulation format used for transmission. Thus, a better match between the bandwidth allocated and the traffic demand can be achieved by properly selecting the type of modulation and the amount of bandwidth. This approach is known as Spectrum-sLICed Elastic optical path network SLICE [5].

In recent years, several research works focused on flexible-grid networks. Many works compared the performance between fixed-grid and flexible-grid networks, focusing mainly on the spectral efficiency and the network blocking probability. It has been proved that flexible-grid networks can improve both metrics [6], [7]. Instead, only few works focused on the energy consumption issue. For instance, in [8] the authors performed a power consumption comparison between flexible and fixed-grid networks. They found that flexible-grid networks can reduce the consumed energy.

However, in most of these works, authors have not devoted particular attention to the design of the IP layer of the network, i.e., the Logical Topology (LT). Indeed, even if flexible-grid networks offer fine spectrum granularity and bandwidth allocation flexibility, accurately designing the logical layer of the network may permit to better exploits these features. Electronic traffic grooming may be performed at the IP layer to reduce the required network interfaces and thus the power consumption. Electronic traffic grooming has been already considered as a strategy to further enhance the spectrum efficiency and the utilization of network resources in flexible-grid networks [9]. However, the traffic grooming must be controlled to avoid a significant increase of power consumption due to Optical/Electronic/Optical (O/E/O) conversions and electronic switching.

In flexible-grid networks, traffic grooming is usually performed using simple greedy heuristics [10] to cope with the network design complexity. The energy-aware design of the LT is not explicitly targeted when traffic grooming is performed in a flexible-grid scenario. Furthermore, an accurate comparison between grooming and no-grooming approach has not been performed. Indeed, heuristics, previously introduced for WDM network design, can not be employed since in the design of flexible-grid networks the additional issues of allocating spectral resources and choosing the suited transmission rate have to be faced. Thus, in this paper we tackle this issue by proposing two greedy heuristics. One heuristic designs the network in an energy efficient way performing electronic traffic grooming, while the other heuristic designs the network without exploiting traffic grooming. We then compare the results for several network scenarios to understand if an energy-aware network design can bring a significant reduction of network energy consumption.

A. Flexible-Grid Network Model

In this work we consider an IP network over an optical flexible-grid network. The physical topology of the network can be represented by a directed graph in which nodes are connected by the physical links existing in the network. At each node of the network, an IP router and a flexible-grid Optical Cross Connect (OXC) are installed. Each physical link

TABLE I.   DETAILS OF THE AVAILABLE MODULATION FORMATS

Modulation Level	BPSK	QPSK	8QAM	16QAM	32QAM	64QAM
Transmission Rate [Gb/s]	12.5	25	37.5	50	62.5	75
Optical Reach [km]	4000	2000	1000	500	250	125

from i to j is characterized by a physical length D_{ij}, expressed in km and such that $D_{ij} = D_{ji}$.

The traffic demands are transmitted from the source to the destination node using lightpaths, which are optical logical channels that can span over one or more physical links. A traffic demand can use one or more consecutive lightpaths to reach the final destination. In this case, the IP router electronically switches the demand between two consecutive lightpaths.

The set of all the established lightpaths forms the LT. Each lightpath is generated at the source node and terminated at the destination node by dedicated flexible OFDM transponders. A flexible OFDM transponder can use any modulation format among the available ones and it is characterized by a maximum transmitting capacity C_{Max} equal to 400 Gb/s. At intermediate nodes the lightpath is transparently switched by the flexible-grid OXC. Since optical switching devices working in a grid-less fashion are not yet available, the spectrum is usually divided in spectrum slots with a much finer granularity than the coarse ITU grid. The optical spectrum on each link is divided in slot of size 12.5 GHz [8], which results in 320 slots per link by dividing the C-band (4 THz) by the slot size. It is also assumed that two empty slots are left as guard-band between two lightpaths so that the OXC can correctly switch the lightpaths.

A given modulation format and a given number of spectrum slots are associated to each lightpath. Each modulation format m is characterized by a maximum bandwidth capacity C_m of a single spectrum slot and by a maximum optical reach in km. The modulation formats considered in this work, their transmission rate and their optical reach are listed in Table I [8]. Depending on the modulation chosen, it is thus possible to create either lightpaths for long distances operating at low bit rate or lightpaths for short distances characterized by very high bit rate. The maximum among the optical reach distances of the available modulation formats corresponds to the maximum reach of the flexible transponder. The maximum number of slots that can be associated to a lightpath with modulation format m is equal to $\lfloor C_{Max}/C_m \rfloor$.

The network design consists in first defining which is the set of lightpaths that can satisfy the traffic demands, i.e., the design of the LT, while optimizing a given design target, for instance the energy consumption minimization. When deciding which lightpaths have to be established, it is required to choose for each lightpath the most suitable modulation and the correct number of slots according to the distance that the lightpath has to cover and the amount of traffic that it has to carry. Finally, slots in the spectrum are assigned to each lightpath, with the constraints that the same set of consecutive slots is assigned to a lightpath over all the physical links that the lightpath is flowing on. Obviously, each slot on a physical link can be assigned only to one lightpath.

The contributions to the network power consumption are given by the IP routers, the flexible OFDM transponders, the flexible OXCs and the Optical Amplifiers (OAs) installed in the physical links. We consider to have available several router configurations, each of them with a given switching capacity of γ Gb/s. The IP router power consumption is computed as $P_{Router} = k \cdot \gamma$, where k is a constant expressed in W/Gb/s that has been set equal to 10 W/Gb/s according to [11] and γ is a multiple of 1 Tb/s. The power consumption values of flexible OFDM transponders, flexible OXCs and OAs have been computed according to the power model reported in [12].

III. Energy-aware Design of Flexible-Grid Networks with Traffic Grooming Heuristic

We propose two heuristics, named the Direct-Lightpath Heuristic (DLH) and the IP-Grooming Heuristic (IGH). The DLH designs the network simply by establishing direct light-paths which can satisfy the traffic demands. The IGH exploits the characteristics of flexible-grid networks to adapt the capacity of already existing lightpaths and thus reducing the number of required network interfaces by exploiting, when convenient, electronic traffic grooming. This involves changing the modulation format and/or the number of spectrum slots used.

The input informations required by both heuristics are the available paths on the physical network topology and the traffic demands. The physical paths are computed using the Dijkstra shortest-path algorithm setting as weights the lengths of the physical links. A set $\mathcal{P}_{ij}$, containing all the paths from i to j ordered for increasing lengths, has been defined for each pair of nodes, i and j belonging to the set of nodes $\mathcal{N}$. The traffic demands, that required to be fulfilled, are organized in the set Λ ordered for decreasing values. If a traffic demand Λ_{sd} is greater than the maximum transmitting capacity C_{Max}, it is split into $\lfloor \Lambda_{sd}/C_{Max} \rfloor$ traffic demands of size C_{Max} and in a traffic demand of size equal to $(\Lambda_{sd} \% C_{Max})$. These traffic demands are allowed to follow different paths from the source to the destination.

The output of both heuristics is the design of the network, denoted with $netDesign$. It consists in the network resources and the network configuration required to satisfy the IP traffic demands. In particular, the design indicates the number of transponders and the router capacity to be installed at each node, and which are the used physical links.

A. The Direct-Lightpath Heuristic

The DLH is described in Algorithm 1. The same operations are repeated for each traffic demand in set Λ (line 1). First, the shortest path from i to j is selected from set $\mathcal{P}_{ij}$ (line 3). Then, until either a feasible path exists (a path is feasible if its length is less than the maximum optical reach of the transmitter) or the demand is not satisfied (line 4), the heuristic tries to establish a lightpath (lines 5-12). The heuristic first selects the best modulation format among the ones available considering the length of the selected path (line 5). The modulation is chosen in such a way that the traffic demand can be entirely allocated in a lightpath using as few as possible spectrum slots. Indeed, a proper choice of the modulation is important because it permits to achieve an efficient use of the spectral

---

Algorithm 1 The Direct-Lightpath Heuristic

Require: $\mathcal{P}_{ij}, \forall (i,j) \in \mathcal{N}, \Lambda$
Ensure: $netDesign$
1: **for all** $\Lambda_{ij} \in \Lambda$ **do**
2:     demand-satisfied $\leftarrow$ **false**;
3:     p=shortest-path($\mathcal{P}_{ij}$);
4:     **while** direct-lightpath-feasible(p) is **true and** demand-satisfied is **false do**
5:         m=select-modulation(Λ_{ij},p);
6:         $slots$=compute-number-required-slots(Λ_{ij},m);
7:         **if** $slots \leq$ available-slots(p) **then**
8:             establish-lightpath(Λ_{ij},p,m);
9:             demand-satisfied $\leftarrow$ **true**;
10:         **else**
11:             p=next-shortest-path($\mathcal{P}_{ij}$);
12:         **end if**
13:     **end while**
14:     **if** demand-satisfied is **false then**
15:         **while** path-exists(p) is **true and** demand-satisfied is **false do**
16:             SP=break-path(p);
17:             subpaths-feasible=check-subpaths(Λ_{ij},SP);
18:             **if** subpath-feasible is **true then**
19:                 establish-subpaths-lightpaths(Λ_{ij},SP);
20:                 demand-satisfied $\leftarrow$ **true**;
21:             **else**
22:                 p=next-shortest-path($\mathcal{P}_{ij}$);
23:             **end if**
24:         **end while**
25:         **if** demand-satisfied is **false then**
26:             **return** No solution is possible;
27:         **end if**
28:     **end if**
29: **end for**
30: slot-assignment=assign-slots($netDesign$);
31: **if** slot-assignment is **true then**
32:     **return** $netDesign$;
33: **else**
34:     **return** No solution is possible;
35: **end if**

---

resources. Thus, a larger amount of traffic can be transmitted in the network with the same network resources. After having chosen the modulation format, the required number of slots is computed (line 6) and the availability of the spectrum slots on the selected path is verified (lines 7-12). If sufficient slots are available, the lightpath is established and the traffic demand allocated to the lightpath (lines 8-9), otherwise the same operations are repeated for the next feasible shortest path in $\mathcal{P}_{ij}$ (line 11).

In the case that the demand is still unsatisfied (line 14), not feasible paths are explored (line 15). The currently selected path is divided in several sub-paths, each of them having a length shorter than the maximum optical reach (line 16). Then, it is verified if each sub-path can support a new lightpath (line 17) by selecting the modulation, computing the required number of slots and verifying the slots availability. If all sub-paths can support the lightpaths (lines 18-20), these are established, otherwise another path is selected if available (line 22). After having explored also all the non feasible paths, if

the traffic demand is still unsatisfied, no solution is possible (lines 25-27) and the algorithm ends. Otherwise, the next traffic demand is selected and the same procedure is repeated.

When all the traffic demands have been fulfilled, spectrum slots are assigned to each lightpath (line 30). If a slot assignment is possible, the solution is validated, otherwise it is rejected. Thus, the slot assignment is performed in a post-processing phase and not during the network design phase.

The heuristic used for the slot assignment is similar to the first-fit assignment algorithm used for wavelength assignment. However, a set of spectrum slots is assigned and not a single wavelength. Furthermore, the slot assignment heuristic has been implemented in such a way that slots are assigned to the lightpaths considering all possible combinations in an exhaustive search. For example, if there are two lightpaths, the heuristic will try to assign slots to lightpath 1 and then to lightpath 2, then it will try first with lightpath 2 and then to lightpath 1. The heuristic terminates if either all combinations have been examined and no feasible solution exists, or a feasible solution is found.

B. The IP-Grooming Heuristic

The main operations performed by IGH are introduced in Algorithm 2. For each traffic demand Λ_{ij} (line 1), IGH first checks if a lightpath from i to j, denoted as l_{ij}, with enough capacity exists (lines 2-3). If it exists, the traffic demand is allocated to this lightpath and a new traffic demand is selected. If no lightpath exists between i and j or there is not enough capacity, the heuristic searches for a sequence of already established lightpaths, called logical path, with starting node i and ending node j (line 5).

In the case that no logical paths from i to j are present, the traffic demand is fulfilled by performing the same operations of one traffic demand of the DLH (lines 6-10). That is, operations from line 2 to line 28 of Algorithm 1 are performed given the current network design and the selected traffic demand. Instead, in the case that at least a logical path exists, the heuristic verifies if it is possible to use this logical path to accommodate the traffic demand (lines 12-21).

Thus, for each available logical path (line 12), it is first checked if enough capacity is available, otherwise the heuristic verifies if it is possible to upgrade the logical path (line 16). The upgrade of a logical paths means to increase the number of spectrum slots and/or to modify the modulation format in the lightpaths not having enough free capacity.

In the case that the traffic can be groomed to the logical path, the power consumption of the adopted solution is computed. If the capacity is already available, transponders are already deployed in the network and, thus, satisfying the traffic demand does not cause an additional power consumption. Only in the case that a router of larger capacity is required, due to the additional traffic to be switched, at some of the intermediate nodes, there would be an additional power consumption. Indeed, the new IP router would present a greater consumption which is taken into account by the heuristic. Instead, if some lightpaths need to be upgraded, there is also an additional power consumption of the transponders due to the lightpaths' upgrade. Indeed, the power consumption of a transponder is

dependent on the modulation format and on the number of slots used, as detailed in Sec. IV-A. When all logical paths have been investigated, the logical path ensuring the smallest increase in the consumed energy is chosen (line 22).

Finally, the power consumption of the "groomed" solution is compared with respect to the case of a "direct-lightpath" solution (line 26). The latter solution is retrieved by satisfying the traffic demand using DLH (line 24) similarly to line 7. The less energy consuming solution is selected (lines 26-33). This procedure is performed for all traffic demands. When all demands have been satisfied, the slot assignment is performed (line 37) and the solution is accepted in case that a feasible slot assignment is possible.

Algorithm 2 The IP-Grooming Heuristic

Require: $\mathcal{P}_{ij}, \forall (i,j) \in \mathcal{N}, \Lambda$
Ensure: $netDesign$
1: **for all** $\Lambda_{ij} \in \Lambda$ **do**
2:    **if** $\Lambda_{ij} \leq$ capacity(l_{ij}) **then**
3:       allocate-demand(Λ_{ij}, l_{ij});
4:    **else**
5:       LP=compute-LP($netDesign, i, j$);
6:       **if** LP is void **then**
7:          $netDesign$=DLH($\Lambda_{ij}, netDesign$);
8:          **if** $netDesign$ is void **then**
9:             **return** No solution is possible;
10:          **end if**
11:       **else**
12:          **for all** $\phi \in LP$ **do**
13:             **if** $\Lambda_{ij} \leq$ capacity(ϕ) **then**
14:                P_ϕ=compute-power(Λ_{ij}, ϕ);
15:             **else**
16:                upgrade=check-upgrade(Λ_{ij}, ϕ);
17:                **if** upgrade is **true then**
18:                   P_ϕ=compute-upgrade-power(Λ_{ij}, ϕ);
19:                **end if**
20:             **end if**
21:          **end for**
22:          groomed-$netDesign$=select-LP($P_\phi, \forall \phi \in LP$);
23:          direct-$netDesign$=DLH($\Lambda_{ij}, netDesign$);
24:          $P_{groomed}$=power-network(groomed-$netDesign$);
25:          P_{direct}=power-network(direct-$netDesign$);
26:          **if** $P_{groomed} \leq P_{direct}$ **then**
27:             $netDesign$=groomed-$netDesign$;
28:          **else**
29:             $netDesign$=direct-$netDesign$;
30:          **end if**
31:          **if** $netDesign$ is void **then**
32:             **return** No solution is possible;
33:          **end if**
34:       **end if**
35:    **end if**
36: **end for**
37: slot-assignment=assign-slots($netDesign$);
38: **if** slot-assignment is **true then**
39:    **return** $netDesign$;
40: **else**
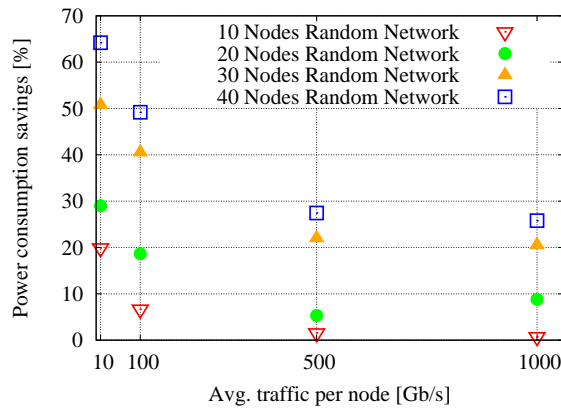41:   **return** No solution is possible;
42: **end if**

Fig. 1. Power consumption savings versus traffic for different sized networks.

Fig. 2. Power consumption savings with confidence interval.

IV. RESULTS

A. Network Scenarios

The heuristics are investigated by performing an exhaustive analysis by changing i) the number of network nodes, ii) the average node degree connectivity, iii) the average link length and iv) the traffic loads. This permits to understand which network's parameters mostly affect the performance of both heuristics and which is their impact on the network power consumption.

Network topologies are chosen by randomly selecting links between pair of nodes until the network is connected and a certain degree of connectivity is ensured. The length of the links is chosen by assigning to the links numbers uniformly generated in the range (0;1), then scaled such that the average link length is the chosen one. More precisely, we consider networks of 10, 20, 30 and 40 nodes. The average node degree connectivity is within the range [3;6]. Several average link lengths are investigated: from 125 km up 2000 km.

Several traffic scenarios are analyzed. Each traffic scenario is identified by the average traffic per node, i.e., a node is generating on average a total amount of traffic. Traffic matrices are created first setting each traffic demands equal to a uniformly generated number in the range (0;1) and then scaling these values to achieve a given target average traffic per node. Four different values of average traffic per node are considered: 10, 100, 500 and 1000 Gb/s.

B. Sensitivity analysis on the number of network nodes

The power savings obtained by IGH over DLH for different sized networks are shown in Fig. 1. In this case the average connectivity has been set to 3 and the average length to 2000 km. Each data point reports the savings for a random network topology and a random traffic matrix. It can be seen that performing traffic grooming leads to significant savings. The savings highly depend on the number of nodes in the network. In larger networks it is indeed possible to achieve higher savings because there are more traffic demands and thus more grooming can be performed. This results in saving a larger number of transponders largely reducing the network power consumption. However, the savings are not linearly increasing with the number of nodes due to the selected IP
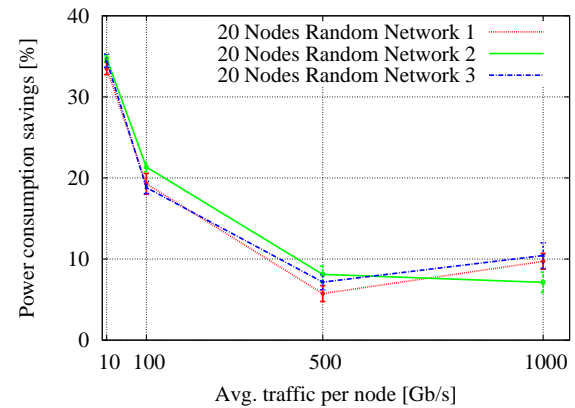
routers. In some scenarios, routers with higher capacity are required causing a reduction of the expected savings.

Savings are larger for low traffic scenario and decrease steadily moving towards high traffic scenarios. Indeed, in high traffic scenarios traffic demands are closer to the maximum lightpath capacity and thus grooming can be rarely performed. Furthermore, it is more convenient to transmit the traffic using a direct lightpath to avoid electronic switching. Indeed, electronic switching should be avoided as much as possible to exploit low capacity IP routers showing a smaller power consumption.

Instead, in low traffic scenarios, large savings can be achieved because it is possible to groom several traffic demands. Indeed, during the network design, it is easier to find already existing logical paths with enough capacity or which can be upgraded to accommodate a new traffic demand, thus resulting in a reduction of the network interfaces.

C. Sensitivity analysis on topologies and traffic matrices

Savings obtained for three different randomly generated networks with 20 nodes are shown in Fig. 2. All networks have connectivity equal to 3 and average link length equal to 2000 km. Confidence intervals, computed using the t-Student distribution over ten different randomly generated traffic matrices, are also shown. The plots show that savings are marginally influenced by the network topologies. A similar observation holds for the savings when changing traffic matrices, as shown by the small confidence intervals.

D. Sensitivity analysis on the average node degree connectivity

Fig. 3 shows the power savings obtained for 20 nodes networks with different connectivity. We consider three different random network topologies for each value of degree connectivity. The traffic load is fixed to 100 Gb/s and the average link length is 2000 km. As the connectivity increases, there is a slight reduction in the advantage obtained by IGH. The reason is that at low connectivity the distance between two nodes is on average much longer than at high connectivity. Thus, at low connectivity the DLH has to set up several consecutive lightpaths to satisfy a traffic demand due to the constrain of the maximum transmission reach of a lightpath. Instead, at high connectivity the distance is on average shorter
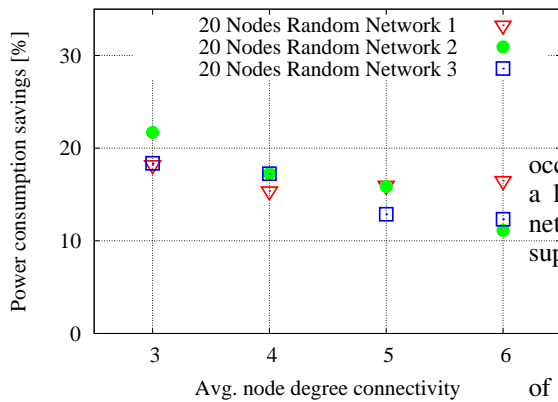
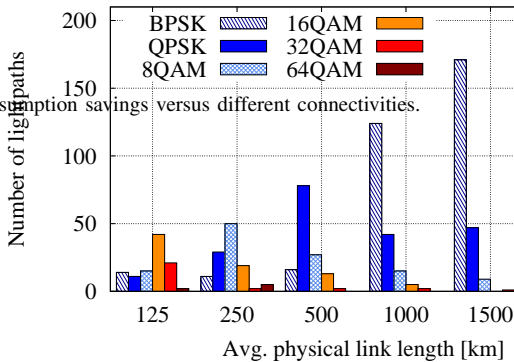Fig. 3.   Power consumption savings versus different connectivities.

Fig. 4.   Modulation formats used versus different link lengths.

occupy a larger number of slots because they transmit with a low bit-rate modulation format. Thus, the capacity that a network with longer links can support is less than the one supported by a network with short links.

V.   CONCLUSIONS

In this paper we address energy-efficiency in the design of flexible-grid networks. This issue has been tackled by designing in an energy-aware manner the network considering both logical and physical layers.

The main contribution of this paper is the evaluation of the importance of performing electronic traffic grooming. We evaluate the grooming impact solving the network design with two greedy heuristics. One heuristic designs the network by simply establishing direct lightpaths between nodes, while the other exploits the bandwidth flexibility to use as much as possible the already installed resources.

Results show that performing electronic traffic grooming allows to achieve very high savings since it permits to effectively exploit the capacity provided by network interfaces, reducing their total number and the associated power consumption. Thus, an optimized design of the logical layer is important also in flexible-grid networks.

REFERENCES

[1] S. . Report, "Enabling the low carbon economy in the information age," http://www.theclimategroup.org, Tech. Rep., 2008.

[2] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, "Energy consumption in optical ip networks," *J. Lightwave Technol.*, vol. 27, no. 13, pp. 2391–2403, Jul 2009.

[3] A. Ahmad, A. Bianco, E. Bonetto, L. Chiaraviglio, and F. Idzikowski, "Energy-aware design of multilayer core networks [invited]," *J. Opt. Commun. Netw.*, vol. 5, no. 10, pp. A127–A143, Oct 2013.

[4] A. Bianco, E. Bonetto, D. Cuda, G. Castillo, M. Mellia, and F. Neri, "Energy-efficient design of wavelength-routing networks," in *Optical Communication (ECOC), 2010 36th European Conference and Exhibition on*, 2010, pp. 1–6.

[5] B. Kozicki, H. Takara, Y. Tsukishima, T. Yoshimatsu, K. Yonenaga, and M. Jinno, "Experimental demonstration of spectrum-sliced elastic optical path network (slice)," *Opt. Express*, vol. 18, no. 21, pp. 22 105–22 118, Oct 2010.

[6] K. Christodoulopoulos, I. Tomkos, and E. A. Varvarigos, "Elastic bandwidth allocation in flexible ofdm-based optical networks," *J. Lightwave Technol.*, vol. 29, no. 9, pp. 1354–1366, May 2011.

[7] K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, "Spectrally/bitrate flexible optical network planning," in *Optical Communication (ECOC), 2010 36th European Conference and Exhibition on*, 2010, pp. 1–3.

[8] J. L. Vizcano, Y. Ye, and I. T. Monroy, "Energy efficiency analysis for flexible-grid ofdm-based optical networks," *Computer Networks*, vol. 56, no. 10, pp. 2400 – 2419, 2012, green communication networks.

[9] S. Zhang, C. U. Martel, and B. Mukherjee, "Dynamic traffic grooming in elastic optical networks." *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 1, pp. 4–12, 2013.

[10] A. Patel, P. Ji, J. Jue, and T. Wang, "Traffic grooming in flexible optical WDM (FWDM) networks," in *OptoeElectronics and Communications Conference (OECC), 2011 16th*, 2011, pp. 405–406.

[11] W. Van Heddeghem, F. Idzikowski, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, "Power consumption modeling in optical multilayer networks," *Photonic Network Communications*, vol. 24, no. 2, pp. 86–102, 2012.

[12] J. L. Vizcaino, Y. Ye, V. López, F. Jiménez, R. Duque, and P. Krummrich, "On the energy efficiency of survivable optical transport networks with flexible-grid," in *European Conference and Exhibition on Optical Communication*. Optical Society of America, 2012, p. P5.05.

than the maximum reach allowing DLH to satisfy the traffic demands using a single lightpath. The IGH instead is marginally influenced by the connectivity because at low connectivity it can reduce the number of network interfaces by performing grooming of traffic demands.

E. Sensitivity analysis on the average link length

Results, not shown due to space constraints, have been retrieved with IGH and DLH for 20 node networks with varying average link length, network connectivity equal to 3 and average traffic per node equal to 100 Gb/s. The power consumption savings are almost not dependent on this parameter and a clear trend can not be noticed. Indeed, increasing the link length increases the power consumption of the physical layer for both heuristics. Thus, the power savings are almost the same for different link length scenarios. A small penalty for the DLH is present for longer links, since it has to establish several consecutive lightpaths instead of a single lightpath in order to satisfy the maximum transmission reach constraint, while the IGH can reduce the power consumption by performing grooming which results in decreasing the network interfaces.

The main effect of link lengths can be noticed in Fig. 4, reporting the distribution of the number of lightpaths with different modulation formats for the solutions retrieved by the IGH. It can be observed that at lower link lengths the utilization of higher modulation formats is more frequent, while as the link length increases the number of lightpaths using lower modulation formats starts to increase. At higher link lengths, the lightpaths, to satisfy the traffic demands, have thus to