Green Horizon: Looking at Backbone Networks in 2020 from the Perspective of Network Operators

Filip Idzikowski*, Luca Chiaraviglio†, Raúl Duque‡, Felipe Jiménez‡ and Esther Le Rouzic§
*Technische Universität Berlin, TKN, Berlin, Germany, filip.idzikowski@tu-berlin.de
†COATI, INRIA, I3S, CNRS, Université de Nice Sophia, Sophia Antipolis, France, luca.chiaraviglio@inria.fr
‡Telefónica Investigación y Desarrollo SA, Madrid, Spain, e.pt1@tid.es, felipej@tid.es
§Orange Labs, Networks and Carriers, Lannion, France, esther.lerouzic@orange.com

***Abstract*—The problem of reducing power consumption in backbone networks by adopting sleep modes has been extensively studied in the literature. In contrast to previous work, we forecast power consumption and traffic in two operator networks (France Telecom (FT) and Telefónica Investigación y Desarrollo (TID)) for the year 2020, and consider both Single Line Rates (SLRs) and Mixed Line Rates (MLRs) deployments. Given these realistic scenarios, the set of devices in sleep mode is computed autonomously on each network node by a simple power-aware algorithm using only local information as input and requiring no rerouting. Results show that the savings obtained by adopting sleep modes are significant, even if power efficiency of network devices (W/Gbps) is increased. Moreover, we show that networks using MLR always consume less energy than SLR-based networks when the power-aware algorithm is applied to the considered scenarios.**

I. Introduction

Power consumption of backbone networks is increasing and becoming a major issue for operators [1]. Although the devices are becoming more and more energy-efficient in terms of power consumption per provided bit rate (W/Gbps) [2], their overall consumption is higher, and the bit rate is not fully utilized [3]. Several approaches aiming at introducing power-awareness into backbone networks have been studied in the literature (see [4] for an overview). One of the most promising approaches is the adoption of sleep modes. When a device of a backbone network is in a sleep mode, the other network devices that remain active have to guarantee Quality of Service (QoS) requirements. While the performance of sleep modes has been widely investigated in the literature, little attention has been paid to the definition of case studies where the effectiveness of sleep modes in real scenarios is evaluated. In fact, most of the works in the literature adopt synthetically generated topologies [5], [6], or case studies which only partially resemble the real networks [7], [8], [9]. Obtaining sensitive information, such as the real Internet Protocol (IP) topology, the power consumption figures and a set of complete traffic matrices for a commercial network is a difficult task. Furthermore, previous works do not focus on the definition of "green horizon" scenarios, i.e., expectations for the forthcoming years.

In this work, we follow a different approach by evaluating the impact of sleep modes on two different scenarios provided by FT and TID. We adopt a simple, fully distributed power-aware algorithm for the selection of the set of devices in sleep mode. To the best of our knowledge, this is the first work that evaluates a power-aware algorithm on evolution scenarios coming from operators. In particular, starting from current measurements of power and traffic, we derive a long-term analysis using forecasted values for devices in 2020. In this way, our work is more tailored to the long-term, making our results representative for future years. Finally, we consider two possible deployment situations. In the first case, we assume that all the installed devices have the same capacity, and we refer to this case as SLR. Then, we consider also the case in which different capacity granularities are allowed, and we refer to it as MLR. The comparison among SLRs and MLRs lets us derive the best choice for the 2020 forecasts, and whether it is worth to maintain the devices of small bit rates in the network in order to save energy by adapting capacity of logical links with fine granularity. Differently from [10] which is focused on network design, we study operation of such SLR and MLR networks with dynamic usage of sleep modes.

The closest paper to our work is [11], in which an impact analysis of power management techniques is performed on the Telecom Italia network (evolution for 2015–2020) and the Greek Research and Technology Network GRNET. In particular, the authors consider the number of installed devices, the power requirements, and the traffic variation as input parameters for a model that evaluates the impact of different power management policies on energy consumption in various network segments. In this work we corroborate the analysis of [11] on two case studies. In particular, we go three steps further by: i) considering the network topology and not solely the number of installed devices, ii) adopting the routing policies used by the operators, iii) assessing the performance of sleep modes over the two scenarios when a specific power-aware algorithm is applied.

The paper is organized as follows. We first present the network model and the power-aware algorithm in Section II. The detailed network scenarios including topologies, traffic patterns and power assumptions are reported in Section III.

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257740 (Network of Excellence "TREND"). This work was carried out during the tenure of an ERCIM "Alain Bensoussan" Fellowship Programme, supported by the Marie Curie Co-funding of Regional, National and International Programmes (COFUND) of the European Commission.

Performance results are presented in Section IV. Finally, Section V concludes our work.

II. Network model and methodology

a) Network model: We consider an IP-over-Wavelength Division Multiplexing (WDM) network, where each node i of a set of nodes V is equipped with an IP router and an Optical Cross-Connect (OXC). The IP routers are interconnected by a set of undirected logical links L. Each logical link consists of one or more lightpaths traversing a set of OXCs interconnected by fibers in the physical layer. A lightpath (a WDM channel which spans one or multiple fiber links) is terminated by transponders and line cards in the IP layer [2]. We assume three types of transponders and line cards depending on their rates, namely 40G, 100G and 400G (G stands for Gbps). Apart from the three SLRs (denoted as slr40G, slr100G and slr400G respectively), we consider also two MLRs. We denote as mlr1 the case in which 100G and 400G devices can be adopted. MLR with all the three available rates is denoted as mlr2.

The network is loaded with a traffic matrix $D(t)$, which is a function of time t. The set of considered time periods is denoted as T. It typically covers 24 hours. Each time period $t \in T$ has duration Δt (usually in the order of minutes). The traffic demands are routed over single paths in the logical topology $(V,\ L)$. $f_l^{ij} \in \{0,1\}$ indicates whether the traffic demand between nodes i and j traverses (independently from t) the logical link $l \in L$. The traffic demands $d^{ij}(t)$ between all nodes $(i,j) \in V \times V$ (stored in the traffic matrix $D(t)$ for the time period $t \in T$) together with their routing f_l^{ij} determine the total traffic $d_l(t)$ flowing through each logical link $l \in L$ at time t:

$$d_l(t) = \sum_{i \in V} \sum_{j \in V} f_l^{ij} \cdot d^{ij}(t) \tag{1}$$

b) Static Base Network: We call the network with all devices constantly active as the Static Base Network (SBN). It is dimensioned for peak traffic matrix D_{SBN}. Each element of D_{SBN} is computed as:

$$d_{SBN}^{ij} = \max_{t \in T} d^{ij}(t) \tag{2}$$

Similarly to Eq. (1), the peak traffic d_l^{SBN} flowing through each logical link $l \in L$ in the SBN is given by:

$$d_l^{SBN} = \sum_{i \in V} \sum_{j \in V} f_l^{ij} \cdot d_{SBN}^{ij} \tag{3}$$

We dimension the SBN depending on the assumed rates. The bit rate C_l of a logical link $l \in L$ is determined by the routing f_l^{ij} of demands d_{SBN}^{ij}. The mapping of each traffic flow d_l^{SBN} into interfaces is arbitrarily made in descending order of bit rates. First, as many highest bit rate interfaces as possible are installed so that their total bit rate does not exceed d_l^{SBN}. Then, this procedure is repeated iteratively for the interfaces of smaller bit rates. Eventually, it is assured in the last iteration that the total bit rate of the logical link is greater or equal to d_l^{SBN}. This results in dimensioning of each logical link $l \in L$ given by:

$$C_l[Gbps] = \left(Y_l^{40G} \cdot 40 + Y_l^{100G} \cdot 100 + Y_l^{400G} \cdot 400\right)[Gbps] \tag{4}$$

where Y_l^{40G}, Y_l^{100G} and Y_l^{400G} denote the number of 40G, 100G and 400G lightpaths active on the logical link $l \in L$.

c) Power-Aware Algorithm: We have adopted the Fixed Upper Fixed Lower (FUFL) algorithm [12] to compute the set of devices in sleep mode at each time period $t \in T$. This algorithm presents features which are particularly appealing from the operator's point of view. It is executed locally on each node using the monitored load on the logical links as input to make decisions. At the same time, FUFL is able not to deteriorate the QoS, while its complexity is low.

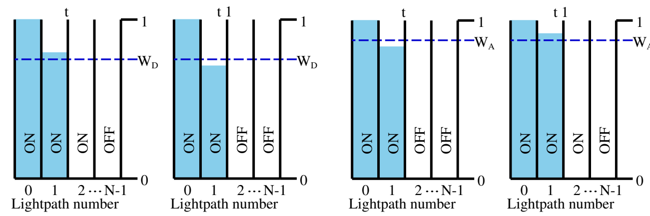

(a) Deactivation (b) Activation

Fig. 1. Utilization of N SLR interfaces on a logical link operated by FUFL

More formally, FUFL is run at each node $i \in V$. It considers the logical links attached to i, and periodically checks the total traffic $d_l(t)$ currently flowing through each of the logical links. The term 'Fixed Upper' corresponds to the fact that the routing f_l^{ij} of demands does not change over time, and the term 'Fixed Lower' corresponds to the fact that the set of lightpaths that can be used in the network and their routing also do not change over time. This means that FUFL needs only local information to make a decision about activation or deactivation of interfaces (line cards and transponders). Fig. 1 shows the main idea of FUFL on a logical link consisting of N parallel lightpaths. Activation and deactivation of a lightpath is triggered by utilization of the previous parallel lightpath, which is determined by $d_l(t)$. The interfaces are deactivated as soon as utilization of the previous parallel lightpath decreases below a pre-defined threshold W_D (Fig. 1(a)). The activation process is triggered by the opposite event, i.e., when utilization of the previous parallel lightpath goes above a pre-defined threshold W_A (Fig. 1(b)). Moreover, in the context of MLR, we performed a preliminary evaluation to check whether power savings are dependent on the order in which network interfaces are deactivated. Whenever high bit rate interfaces are deactivated first, larger power savings are obtained, since deactivation of (finite number of) low bit rate interfaces at the end of the algorithm allows adjustment of the capacity of a logical link with finer granularity.

In contrast to [12], we introduced thresholds W_A and W_D, which can be adjusted by the network operator to account for the QoS and the time needed to activate network devices. Moreover, we allow to shift traffic between all lightpaths on

TABLE I
REFERENCE NETWORKS (LOGICAL TOPOLOGIES) AND TOTAL TRAFFIC DEMAND

Network	Nodes	Links	Nodal degree (Min/Avg/Max)	Link length [km] (working & protection) (Min/Avg/Max)	Total demand [Tbps] Low (LA)	Medium (MA)	High (HA)
France Telecom (FT)	38	72	1/3.79/14	2/270/1327	20	33	50
Telefónica Investigación y Desarrollo (TID)	113	254	1/2.25/30	1/201/2431	39	55	77

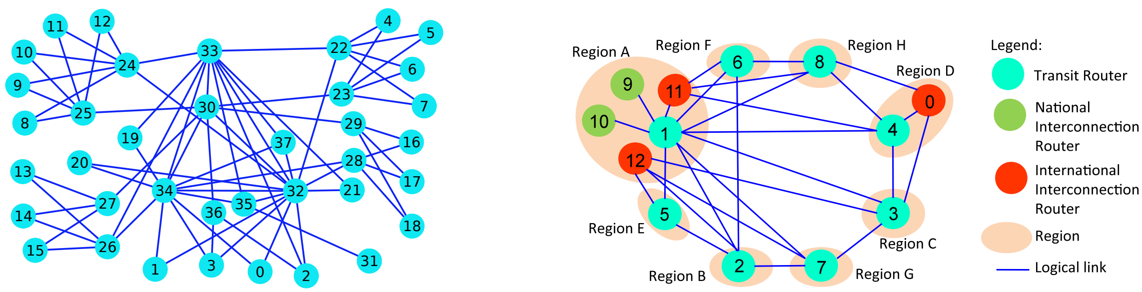

(a) FT

(b) TID (the edge routers in the 8 regions are not shown)

Fig. 2. Logical topologies

a logical link, regardless of their physical realization. However, we neglect the time that is necessary for the traffic shifting between lightpaths of different rates. The traffic shifting does not influence the network stability because the logical topology is unchanged. However neglecting the time needed for the traffic shifting slightly underestimates the energy consumption because of the make-before-break mechanism. Eventually, FUFL application between different network domains is out of scope of our work because there is usually Optical-Electrical-Optical conversion on the border between the domains.

III. Network Scenarios

We apply FUFL to two operator networks, which are representative backbones (realistic, but not real) of FT and TID. Table I reports a summary of the main network features.

Topology The IP (logical) topologies are shown in Fig. 2. The logical topologies are realized over photonic meshed infrastructures with WDM equipment. Both networks are comparable in terms of nodal degree and average link length (see Table I). The link lengths range between few kilometers and a few thousands of kilometers. The shortest ones are adopted to connect nodes close to each other, e.g., in the same building, while the longest ones correspond to submarine links to overseas territories. Finally, we can observe that many links are deployed from the periphery to the center of the topology (region A for TID and nodes 32-34 for FT), where most of the peering points to the Internet are located.

Architecture Both the FT and TID networks have a hierarchical structure composed of three levels: edge, transit and Internet/peering points. Edge routers are the first (lowest) level (not shown in Fig. 2(b) for the sake of readability). Each edge router is connected to a pair of transit routers, which constitute the second level. The interconnection among transit routers and to the Internet gateways and national peering points is done in the third hierarchical level. Most of the nodes have dual router connectivity allowing load balancing. The routing schemes were provided by the operators as part of the scenario [13]. Interfaces between the photonic and IP layers are realized following International Telecommunication Union - Telecommunication Standardization Sector (ITU-T) G.692 recommendation. In particular, line cards with "gray" interfaces are used. They are connected to transponders which terminate "colored" paths across the optical network.

Traffic The complete IP traffic matrix D_{SBN} is provided for both networks for a 2020 forecast with three growth assumptions (Low, Medium, and High Assumption – LA, MA and HA respectively) compared to today's situation (see Table I and [13]). In both networks most of the traffic is massively originated from/to the international interconnections.

The daily variation of the total demand $\sum_{(i,j)\in V\times V} d^{ij}(t)$ is shown in Fig. 3 for two representative days T (working and weekend) and three growth assumptions (LA, MA, and HA). Traffic matrices are provided with a granularity Δt equal to 5 min. A clear day-night pattern is visible for both networks, with traffic variation between night and day reaching almost a factor of 10. This confirms the high network overdimensioning during the night. Moreover, the weekend traffic is consistently lower than the traffic during the working day, suggesting that there is ample room to apply power-aware mechanisms to the two networks.

By investigating the traffic generated by each node over time, we have found out that most of edge nodes present a similar behavior, i.e., almost the same day-night pattern. This is due to two main reasons: aggregation of traffic from/to tens of thousands of users, and the fact that all the nodes are located in the same time-zone.

Power model We focus on transponders and line cards, since these devices are targeted to be powered off by FUFL. In particular, we propose two forecasts of power consumption values for the year 2020, one that we call Capital Expenditures

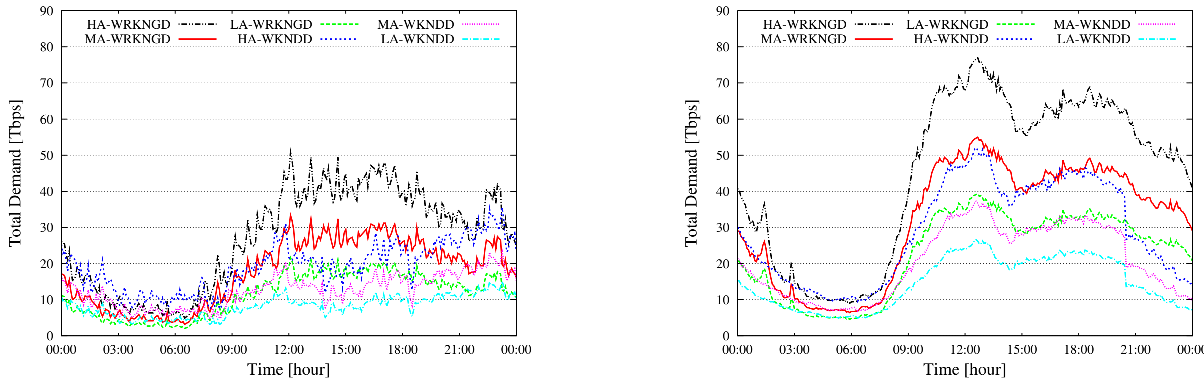

(a) FT

(b) TID

Fig. 3. Daily traffic variation (downlink) for Low/Medium/High traffic Assumption (LA/MA/HA) on working and weekend days (WRKNGD and WKNDD)

TABLE II
POWER CONSUMPTION IN WATT (POWER EFFICIENCY IN WATT/GBPS)

(a) CapEx-driven forecast for 2020

Bit rate [Gbps]	Transponder	Line card	Total
40	100 (2.5)	300 (7.5)	400 (10)
100	120 (1.2)	420 (4.2)	540 (5.4)
400	300 (0.75)	1180 (2.95)	1480 (3.7)

(b) Energy-driven forecast for 2020

Bit rate [Gbps]	Transponder	Line card	Total
40	40 (1)	200 (5)	240 (6)
100	120 (1.2)	420 (4.2)	540 (5.4)
400	350 (0.875)	1250 (3.125)	1600 (4)

(CapEx)-driven and the other that we call Energy-driven. The actual power consumption and power efficiency values for both forecasts are shown in Table II. Note that the power efficiency values (in W/Gbps) account for the bit rate of the devices, and not the actual throughput. We keep the values from [2] as an upper bound for our forecasts. Finally, the power model of transponders is based on the new generation of coherent detection transponders [14] considering large hardware improvement especially for the electronic digital processing part with respect to today's power consumption.

In the CapEx-driven forecast, we assume that the use of 40G interfaces will be marginal resulting in a reduced effort for the development and optimization of 40 Gbps boards that will consume almost the same amount of power as 100 Gbps boards. We base this assumption on the observation that flexible grid networks and variable bit rate transponders get most traction from the industry and that 40G could start becoming "legacy". We thus suppose two different implementations for 40G and 100G i.e., Dual-Polarization Binary Phase-Shift Keying (DP-BPSK) and Dual-Polarization Quadrature Phase-Shift Keying (DP-QPSK), respectively. As they have a common technological basis in this scenario, power consumption of 40G transponders is close to power consumption of 100G transponders. As 100G is more probably the next bit rate for routers, 40G line cards are much less power-efficient compared to 100G line cards. Finally, the 400G interfaces will benefit from large development effort.

In the Energy-driven forecast we assume that for "green" reasons, interfaces of small bit rates will be favored because fine granular interfaces allow for a fine tuning of the bit rate of the logical links as load varies. As a result, 40G interfaces benefit from a large effort to reduce their power consumption, while high bit rate interfaces at 400 Gbps do not. In this scenario, the same modulation format DP-QPSK is used for 40G and 100G. This could be somewhat more expensive for 40G transponders (a single spare part needed), but it offers operational advantages, and better power efficiencies. In particular, the symbol rate of 40G transponders is less than a half of the symbol rate of 100G transponders, and thus the Digital Signal Processor power consumption (which dominates the transponder power consumption) could be significantly reduced. As a result, power efficiency (W/Gbps) of 40G transponder would be improved compared to 100G transponders, while line card power efficiency would be close for 40G and 100G bit rates.

Finally, we assume negligible power consumption in sleep mode for transponders and line cards in both forecasts.

IV. RESULTS

We evaluate FUFL on the FT and TID scenarios, assuming medium traffic assumption MA unless specified otherwise. We set $W_D = 1.0$ and $W_A = 1.0$, which corresponds to deactivation of a lightpath with corresponding devices as soon as it becomes empty, and its activation as soon as the previous parallel lightpath becomes full. Fig. 4 reports the variation of power over time considering a working day, CapEx-driven forecast, MLR and SLR devices. The figure reports also the power consumption of the SBNs, which correspond to the networks with all the devices powered on. Interestingly, FUFL greatly reduces the power consumption of active interfaces in the low demand hours, with a reduction of power of up to 76% for both the FT and TID scenarios. Moreover, while the

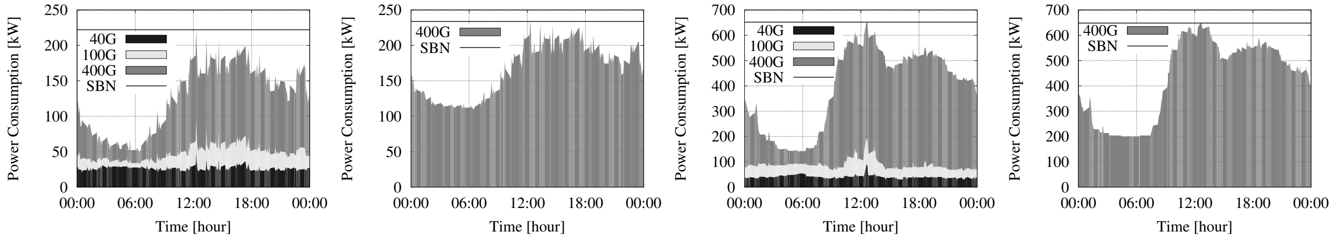

(a) FT, mlr2 (b) FT, slr400G (c) TID, mlr2 (d) TID, slr400G

Fig. 4. Power consumption vs. time for the two networks and different line rate assumptions (working day, medium traffic assumption, CapEx-driven forecast)

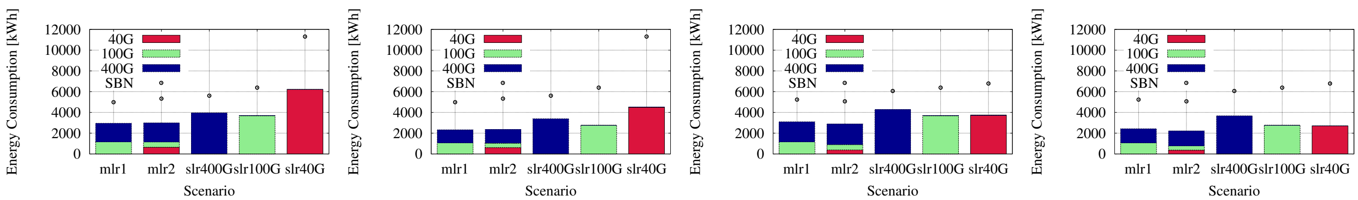

(a) Working Day, CapEx-driven (b) Weekend Day, CapEx-driven (c) Working Day, Energy-driven (d) Weekend Day, Energy-driven

Fig. 5. Daily energy consumption over the FT network (SBN and with FUFL) for different days and power forecasts (medium traffic assumption)

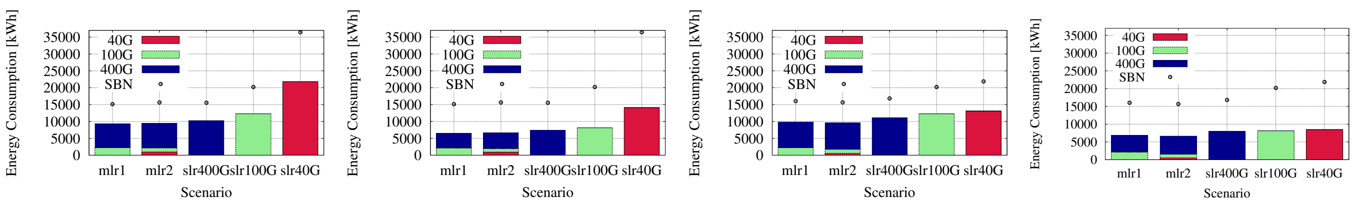

(a) Working Day, CapEx-driven (b) Weekend Day, CapEx-driven (c) Working Day, Energy-driven (d) Weekend Day, Energy-driven

Fig. 6. Daily energy consumption over the TID network (SBN and with FUFL) for different days and power forecasts (medium traffic assumption)

power consumption of the SBN is comparable when using mlr2 and slr400G, the power consumption of the network running FUFL is almost always lower when MLR devices are deployed than the case in which SLR devices are used. E.g., during night hours (2:00 am to 7:00 am), the power consumption of the FT network running FUFL with mlr2 is always lower than 78 kW, while the same network adopting SLR devices consumes minimally 109.52 kW in the same period. This is due to the fact that in the SLR scenario, the high bit rate 400G devices are kept active for connectivity reasons, even if their utilization is low. Finally, note that the power savings in the network adopting MLR are mostly due to deactivation of the 400G interfaces, since these devices are the most power consuming, but at the same time the most power-efficient and thus widely deployed ones in our study. The other types of interfaces are used to fine adjust the bit rate of the logical links to the current load.

In order to have a broader view on the savings that can be brought by FUFL, we looked at the energy consumed by all interfaces in the network over the whole day in both the FT network (Fig. 5), and the TID network (Fig. 6). Focusing first on the SBNs (the same during a working day and a weekend day), we can notice that in general the networks using MLRs consume less energy than the networks using SLRs, although slr400G consumes comparable amount of energy as mlr2 in the TID network. The energy consumption in the CapEx-driven forecast is even slightly lower in the TID's slr400G SBN than in the mlr2 SBN (see Figures 6(a) and (b)). Apart from the high power efficiency of the 400G interfaces in the CapEx-driven forecast, this is due to the fact that the utilization of the last 400G interfaces on logical links is relatively high in the slr400G TID network, namely 44/49/47% for the low/medium/high traffic assumptions respectively. The corresponding values in the FT network do not exceed 1%.

The two MLRs perform comparably both in the SBNs and in the networks running FUFL. mlr1 is better than mlr2 in terms of energy consumption assuming the CapEx-driven forecast, while mlr2 is better than mlr1 assuming the Energy-driven forecast. This is due to the fact that 40G interfaces are used only in mlr1. Their power consumption is high in the CapEx-driven forecast, and relatively low in the Energy-driven forecast (see Table II). The same effect can be observed in the networks using FUFL. However, differently than in the SBNs, the advantage of using the most power-efficient (in W/Gbps) interfaces (400G) cannot be always observed in the networks using FUFL. This is especially noticeable in the FT network, where slr400G is not the least energy-consuming SLR network even for the CapEx-driven forecast of power consumption

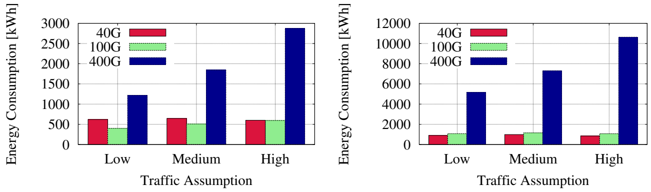

(a) FT Network

(b) TID Network

Fig. 7. Daily energy consumption breakdown over different rates in mlr2 and different traffic assumptions for working days (CapEx-driven forecast)

values, which favors high bit rate interfaces (Fig. 5(a) and (b)).

The power model significantly influences the energy consumed in the networks using FUFL, which is especially reflected by the energy consumption of the 40G interfaces (compare CapEx-driven and Energy-driven forecasts in Figures 5 and 6). Looking at the two MLRs, there is however no significant advantage when using 40G interfaces, even when assuming the Energy-driven forecast. Both MLRs in turn outperform the SLR with 400G interfaces, which are expected to be commonly used in 2020.

Eventually, we note that energy consumption during a weekend day is lower than during a working day over all considered scenarios. The average energy saving during a weekend day with respect to a working day over all five rate assumptions equals (CapEx-driven forecast/Energy-driven forecast) 891.18/778.08 kWh in the FT network, and 4066.55/3577.54 kWh in the TID network.

Network operators need to make long-term forecasts of the traffic in order to plan operation of their networks. In Fig. 7 we consider the low/medium/high traffic assumptions on the working days in the FT and TID networks using all the three types of network interfaces (mlr2) with CapEx-driven forecast. The daily energy consumption consistently grows with the assumed traffic only for the biggest network interfaces. Power consumption of 40G and 100G interfaces remains comparable over all the traffic assumptions. This is a result of the chosen rule to use as many 400G interfaces as possible in the SBN (see Section II). Consequently, the 400G interfaces carry the majority of the traffic in the SBN, even though their absolute count is the highest only for the high traffic assumptions in both networks, and medium assumption in the TID network. 40G interfaces consume more energy than 100G interfaces in the FT network, while the opposite holds for the TID network. This can be explained by the number of installed interfaces in both networks. The lower the traffic, the higher percentage of interfaces of small bit rate are installed. In particular, the 40G interfaces account for 36/45/52% of all the installed interfaces in the FT network for high/medium/low traffic assumption respectively. Corresponding values for the TID network are lower, namely 25/31/36%.

V. Conclusion & Future Work

We have evaluated the performance of a fully distributed algorithm called FUFL in terms of saving energy in the TID and FT networks with traffic and power forecasts for the year 2020. The savings obtained by adopting FUFL are significant, with up to 76% of saved power compared to an always on solution, showing that there is ample room in pursuing energy-efficient approaches even for future networks. The adoption of MLRs provides benefits over SLRs in all considered scenarios, when FUFL is applied. However, there is no clear indication that interfaces of small bit rates should be developed in order to provide the flexibility of adaptation to changing loads in the network. Using both power consumption forecasts considered in our study, a network with MLR of all three types of interfaces showed only slightly lower energy consumption than the network with MLR of only 100G and 400G interfaces. In addition, maintenance of devices of different technologies may be difficult for network operators, especially in the context that old equipment may be more prone to failures. The disposal of old equipment is left for future work. Moreover, we plan to implement FUFL on a testbed checking its sensitivity to W_D and W_A.

References

[1] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch, "Energy Consumption of Telecommunication Networks and Related Improvement Options," *Journal of Selected Topics in Quantum Electronics*, vol. 17, no. 2, pp. 285–295, March-April 2011.

[2] 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, October 2012.

[3] A. Odlyzko, "Data networks are lightly utilized, and will stay that way," *Review of Network Economics*, vol. 2, no. 3, pp. 210–237, September 2003.

[4] R. Bolla, R. Bruschi, F. Davoli, and F. Cucchietti, "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 2, pp. 223–244, 2011.

[5] L. Chiaraviglio, D. Ciullo, E. Leonardi, and M. Mellia, "How much can the internet be greened?" in *Proc. of the GLOBECOM Workshop on Green Communications, Honolulu, USA*, December 2009.

[6] F. Musumeci, M. Tornatore, R. Fiandra, and A. Pattavina, "The role of network topology on the energy efficiency of ip-over-wdm architectures," in *Proc. of the ONDM, Colchester, UK*, 2012.

[7] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," in *Proc. of the INFOCOM, Phoenix, USA*, April 2008.

[8] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, "An energy saving routing algorithm for a green OSPF protocol," in *Proc. of the INFOCOM Workshops, San Diego, USA*, March 2010.

[9] L. Chiaraviglio, M. Mellia, and F. Neri, "Energy-aware backbone networks: a case study," in *Proc. of the ICC Workshop on Green Communications greencom, Dresden, Germany*, June 2009.

[10] P. Chowdhury, M. Tornatore, and B. Mukherjee, "On the Design of Energy-Efficient Mixed-Line-Rate (MLR) Optical Networks," *Journal of Lightwave Technology*, vol. 30, no. 1, pp. 130–139, January 2012.

[11] R. Bolla, R. Bruschi, A. Carrega, F. Davoli, D. Suino, C. Vassilakis, and A. Zafeiropoulos, "Cutting the energy bills of Internet Service Providers and telecoms through power management: An impact analysis," *Computer Networks*, vol. 56, no. 10, pp. 2320–2342, July 2012.

[12] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz, "Dynamic routing at different layers in IP-over-WDM networks – maximizing energy savings," *Optical Switching and Networking, Special Issue on Green Communications*, vol. 8, no. 3, pp. 181–200, July 2011.

[13] O. Renais and J. L. Briand, "Tracks for transport network architecture optimization," in *Proc. of the Networks, Rome, Italy*, October 2012.

[14] A. Morea, S. Spadaro, O. Rival, J. Perellz, F. A. Bujan, and D. Verchere, "Power management of optoelectronic interfaces for dynamic optical networks," in *Proc. of the ECOC, Geneva, Switzerland*, September 2011.