

Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios

Filip Idzikowski Sebastian Orłowski Christian Raack Hagen Woesner, Adam Wolisz
TKN, TU Berlin, Germany atesio GmbH, Berlin, Germany ZIB, Berlin, Germany TKN, TU Berlin, Germany
idzikowski@tkn.tu-berlin.de orlowski@atesio.de raack@zib.de {woesner,wolisz}@tkn.tu-berlin.de

Abstract—We estimate potential energy savings in IP-over-WDM networks achieved by switching off router line cards in low-demand hours. We compare three approaches to react on dynamics in the IP traffic over time, FUFL, DUFL and DUDL. They provide different levels of freedom in adjusting the routing of lightpaths in the WDM layer and the routing of demands in the IP layer. Using MILP models based on realistic network topologies and node architectures as well as realistic demands, power, and cost values, we show that already a simple monitoring of the lightpath utilization in order to deactivate empty line cards (FUFL) brings substantial benefits. The most significant savings, however, are achieved by rerouting traffic in the IP layer (DUFL), which allows emptying and deactivating lightpaths together with the corresponding line cards. A sophisticated reoptimization of the virtual topologies and the routing in the optical domain for every demand scenario (DUDL) yields nearly no additional profits in the considered networks.

I. INTRODUCTION

In the light of scarce fuel resources and the rising demand for energy [1], there is a growing interest in solutions and “green” strategies in different fields to reduce the power consumption. In this work, we focus on energy efficiency in IP (Internet protocol) over WDM (wavelength division multiplexing) backbone networks. The conservative approach to network design leads to static solutions covering all the traffic patterns and minimizing capex (capital expenditures). A dynamic adaptation of the network topology and the number of active components to the traffic patterns influences opex (operational expenditures), where energy is one of the key factors. The power consumption of IP routers and line cards is almost independent of the load and reaches hundreds of kW for multi-shelf configurations [2, 3, 4]. We pose the question of how much energy can be saved by dynamically switching off idle IP router line cards in low-demand hours. The aim of this paper is to estimate and compare the potential energy savings of three different approaches to make line cards idle by reconfiguring the routing at the IP and/or WDM layer.

Although several papers have focused on power consumption in single- [3, 5, 6, 7] and multi-layer networks [8, 9, 10], our work is, to the best of our knowledge, the first study comparing the contribution of rerouting at different layers to the energy savings. We systematically investigate the influence of traffic variability on power consumption of dynamically reconfigurable networks using realistic data on dynamic traffic, network topologies, costs, and power of single network elements. As shown in Section III, we use variations of the

same mixed integer linear program (MILP) to design a static base network and to reconfigure the network in every demand scenario such as to maximize the number of idle line cards. This approach allows us to provide energy-optimal solutions or at least upper bounds on the potential savings. Section IV describes the used data. The computational study in Section V reveals that allowing dynamic routing at the IP layer depending on the traffic pattern contributes the most to the energy savings. Reconfiguring lightpaths in the WDM layer gives only little additional benefit. Section VI concludes our work.

II. NETWORK MODEL

We focus on IP-over-WDM networks, where the WDM layer offers optical by-pass technology. Nodes in the WDM layer, which represent optical cross-connects (OXC), are interconnected by links representing optical fibers. Each fiber carries up to B WDM channels of capacity C Gbps each. OXCs may connect incoming WDM channels to outgoing ones (assuming full wavelength conversion capability), or terminate them in the corresponding nodes in the IP layer. The IP layer is interconnected with the WDM layer by colored router line cards, which perform optical-electrical-optical (OEO) conversion. IP routers can be equipped with line cards of bandwidth C Gbps. Lightpaths, which are concatenations of WDM channels, terminate in the line cards. All parallel lightpaths between two IP routers form a virtual link in the IP layer. The virtual links together with the IP routers form a virtual topology. IP routers are sources and destinations of backbone traffic, which is converted into an optical signal by the line cards and directly fed into OXCs.

A lightpath between two particular line cards may be realized by one or several physical paths (routing in the WDM layer). The IP traffic demands are routed on the virtual links defined by the set of lightpaths. We allow multi-path routing both in the IP layer (IP traffic can be split and routed via multiple virtual paths) and in the WDM layer (more than one lightpath may be established between two IP routers).

A (capex) cost-minimized static multi-layer network serves as a starting point to our investigations. Given demands with temporal and spatial dynamics, it is designed to accommodate all traffic without any adaptation of the routing and hardware configuration. Based on this static base network, we consider three different approaches to decrease power consumption in the operational phase by switching off unused line cards.

Fixed Upper Fixed Lower (FUFL): Both the routing of IP traffic in the upper virtual layer and the realization of lightpaths in the lower WDM layer are fixed over time. Demands have to be routed as in the static base network, using the same lightpaths with the same percental splitting as in the base network. Line cards of empty lightpaths can be switched off. Traffic can be shifted only between parallel lightpaths that correspond to the same physical path.

Dynamic Upper Fixed Lower (DUFL): The virtual topology (including the realization of lightpaths) is fixed as in FUFL (Fixed Upper), but the routing of IP traffic can be changed (Dynamic Upper). In every demand scenario, we aim to route the IP demands in the virtual topology in such a way that as many lightpaths as possible are emptied, which allows to switch off the corresponding line cards.

Dynamic Upper Dynamic Lower (DUDL): Both the routing of the IP traffic in the virtual layer and the realization of lightpaths in the physical layer can be changed over time, with the only restriction that the number of line installed cards at each IP router must not be exceeded. The number of used line cards has to be minimized by jointly optimizing the routing in the IP and WDM layer.

We do not consider the Fixed Upper Dynamic Lower approach, since the IP routing would not be able to react to the virtual topology change in this case.

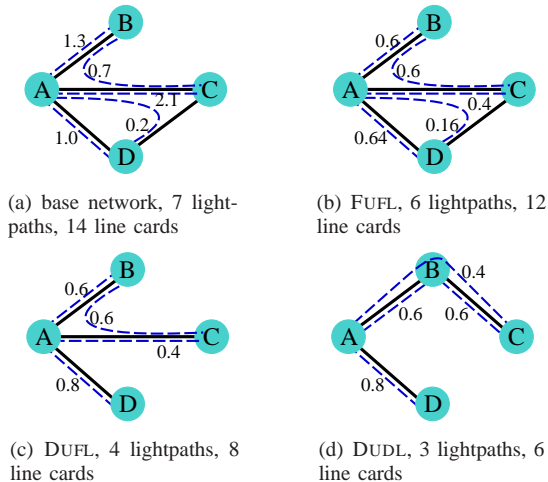


Figure 1. Example of optimal virtual topologies (solid lines) and IP routings (dashed lines). Granularity of virtual link capacity is $C = 1$ (two line-cards). There are four peak demands ($AB = 1.3$, $AC = 2.1$, $AD = 1.2$, $BC = 0.7$), which decrease in the low-demand hour ($AB = 0.6$, $AC = 0.4$, $AD = 0.8$, $BC = 0.6$). We save 2, 6, and 8 line cards with FUFL, DUFL, and DUDL, respectively.

Fig. 1 shows a simple example on how each of the approaches may decrease the number of active line cards in a low-demand hour. The physical fiber installation and the hardware configurations at the nodes from the base network are untouched. New line cards must not be installed.

FUFL is the most restrictive option. It is the easiest to be realized in practice since it does not require any optimization but only monitoring of the lightpath utilization. Decisions on switching line cards on and off can be taken locally. Its

drawback is to rely on the routing defined by the base network, which can be suboptimal especially in low-demand hours.

In contrast, DUFL and DUDL with the objective of minimizing lightpaths are NP-hard optimization problems, as they generalize the uncapacitated fixed charge flow problem [11]. DUFL is a single-layer network design problem which can be solved to optimality in a reasonable amount of time in practice, see Section V. DUDL is a computational challenge since it involves optimizing two coupled network layers simultaneously, similar to designing the base network.

Dynamics in the IP routing (DUFL) may allow more line cards to be switched off, compared to FUFL, by choosing a smart IP routing in each demand scenario, but it may lead to instabilities of connection-oriented protocols (e.g. due to overtaking of packets upon the change of the IP routing). Moreover, decisions about the IP routing changes need to be forwarded to all the involved routers. Even more signaling is needed for dynamics in the WDM layer (DUDL). It has to be ensured that no packets are lost in the reconfiguration phase, when lightpaths are torn down.

The detailed study of such operational issues as mentioned above is beyond the scope of this paper. We also do not provide an algorithm or protocol to actually reconfigure the network when the demand changes. Indeed, the goal of this paper is to compare the three approaches from a conceptual perspective and to give an upper bound on their energy-saving potential. In this respect, the savings with DUDL serve as an upper bound for those with DUFL, which in turn serves as a benchmark for the more restrictive FUFL. In practice, a trade-off between potential energy savings and the complexity of reconfiguration needs to be found on a given time scale.

III. METHODOLOGY AND MATHEMATICAL MODELS

We first design an IP-over-WDM network that serves as a basis for our investigations. This base network is considered to be *static*: independent of demand fluctuations over time, all hardware equipment as well as the IP-routing and the realization of lightpaths in the optical domain are fixed. The network is designed to accommodate all (peak) demand realizations in a given period of time without any reconfiguration. Among all possible topologies and routings, we aim at starting with a base network that is minimal with respect to equipment cost. Given a base network, we then compare the three approaches to minimize the number of active line cards.

A. Design of a base network

Designing a cost-minimal network that allows all known demand scenarios to be realized is a highly complex robust network design problem, for which only first mathematical approaches have been developed (see [12] and references therein). Our approach is based on constructing a demand matrix that refers to all peak demands over time and dimensioning the base network with respect to this maximum matrix.

Let V be the set of all demand end-nodes and let d_{ij}^t be the (undirected) demand value for each pair of nodes $(i, j) \in$

$V \times V, i < j$ and each point in time $t \in T$. We compute a maximum demand matrix $(d_{ij}^{\max})_{(i,j) \in V \times V}$ by

$$d_{ij}^{\max} := \max_{t \in T} d_{ij}^t, \quad (1)$$

and calculate a minimum-cost IP-over-WDM network which satisfies this maximum demand matrix.

We optimize both network layers at the same time in an integrated step. Our model is close to the one used in [13]. It comprises all relevant sources of installation cost both in the IP and the WDM layer. Extensions of this model are later used to evaluate the energy savings in different demand scenarios.

Parameters Assuming all network elements to be bidirectional, we model the optical layer by an undirected physical supply network $G = (V, E)$ consisting of the locations V and the physical links E . Every node $i \in V$ has a preinstalled optical cross-connect and can be equipped with an IP router out of the set N of IP routers. Every router $n \in N$ has a maximum switching capacity of C^n and a cost of α^n . Every physical link $e \in E$ can operate an arbitrary number of fibers at cost β^e per fiber, each supporting B wavelength channels.

For every node pair $(i, j) \in V \times V, i < j$, the set $P_{(i,j)}$ denotes all admissible routing paths in G between nodes i and j , which can be used to realize lightpaths. Let P be the union of all these paths and P_i the set of all paths ending at node i . Each path $p \in P$ can be equipped with multiples of the bandwidth C . This corresponds to installing an appropriate number of interface cards at the end-nodes of p . The cost γ of establishing a lightpath (installing one unit of capacity C on a path $p \in P$) is given by the cost for two line cards. Every such unit consumes one wavelength channel in the physical network on every physical link of the path.

Demands and commodities From the matrix d^{\max} , we construct commodities $k \in K$ by aggregating demands at a common source node. This modeling trick [13] reduces the number of commodities from $\mathcal{O}(|V|^2)$ to $\mathcal{O}(|V|)$ and leads to commodities having one source and several target nodes. A net demand value d_i^k is associated with every commodity $k \in K$ and every node $i \in V$ such that $\sum_{i \in V} d_i^k = 0$. It specifies the net demand starting (if $d_i^k \geq 0$) or ending (if $d_i^k < 0$) at node i . The demand value of a commodity is given by $d_i := \sum_{k \in K} |d_i^k|$.

Variables The flow variables $f_{ij}^k, f_{ji}^k \in \mathbb{R}_+$ describe the flow for commodity k on virtual link between i and j in both directions. These variables aggregate the IP traffic on all lightpaths with endnodes i and j . The distribution of virtual link flow to the chosen physical representations and also the disaggregation of commodities to individual demands can be done in a post-processing step, see [14]. Variables $y_p \in \mathbb{Z}_+$ count the number of lightpaths realized on $p \in P$. Similarly, $y_e \in \mathbb{Z}_+$ denotes the number of fibers installed on physical link $e \in E$. The binary variable $x_i^n \in \{0, 1\}$ states whether or not router n is installed at node $i \in V$.

Model The problem of minimizing the cost for a feasible network configuration and routing satisfying the demand matrix d^{\max} can be formulated as the MILP (2). Equations (2a) are the flow conservation constraints for every node

and commodity. Inequalities (2b) choose a subset of paths between the nodes i and j and install enough capacity to accommodate all the virtual link flow corresponding to (i, j) . The virtual node capacity constraints (2c) make sure that the capacity of a node suffices to switch all the incoming traffic, including the emanating demand. Constraints (2d) select one router configuration at every node. Eventually, the physical link capacity constraints (2e) make sure that the number of available wavelengths on a fiber is not exceeded.

$$\min \sum_{i \in V, n \in N} \alpha^n x_i^n + \gamma \sum_{p \in P} y_p + \sum_{e \in E} \beta^e y_e$$

$$\sum_{j \in V \setminus \{i\}} (f_{ij}^k - f_{ji}^k) = d_i^k, \quad i \in V, k \in K \quad (2a)$$

$$\sum_{p \in P_{(i,j)}} C y_p - \sum_{k \in K} (f_{ij}^k + f_{ji}^k) \geq 0, \quad (i, j) \in V \times V \quad (2b)$$

$$\sum_{n \in N} C^n x_i^n - \sum_{p \in P_i} C y_p \geq d_i, \quad i \in V \quad (2c)$$

$$\sum_{n \in N} x_i^n \leq 1, \quad i \in V \quad (2d)$$

$$B y_e - \sum_{p \in P: e \in p} y_p \geq 0, \quad e \in E \quad (2e)$$

$$f_{ij}^k, f_{ji}^k \in \mathbb{R}_+, y_p, y_e \in \mathbb{Z}_+, x_i^n \in \{0, 1\} \quad (2f)$$

B. Evaluation of different demand scenarios

In the following, we explain how we adapt model (2) to evaluate the possible energy savings for FUFL, DUFL and DUDL under dynamically changing demands.

FUFL Consider a demand between nodes i and j with base value d_{ij}^{\max} . In each low-demand hour $t \in T$ where this demand has value d_{ij}^t , we reduce the flow on each virtual path for this demand by the factor $d_{ij}^t / d_{ij}^{\max} \in [0, 1]$ and reduce the capacity on the path accordingly. No optimization is needed.

DUFL For every $t \in T$, we compute a new routing by using a variant of (2). We fix the physical network and the IP routers by fixing the variables y_e and x_i^n to their values from the static base network. The virtual link capacity variables y_p can be reduced with DUFL compared to the base network, but not augmented. This is ensured by adding the constraint $y_p \leq y_p^{\text{base}}$ to (2) for all $p \in P$, where y_p^{base} is the capacity of virtual link p in the base network. The IP flow can be rerouted such as to minimize the number of active line cards. Hence an energy optimal routing for every time period $t \in T$ can be computed by fixing and bounding variables y_e, x_i^n and y_p in (2) as described above, using the demand matrix d^t instead of d^{\max} , and changing the objective function to minimize $\sum_{p \in P} y_p$.

DUDL Similarly to DUFL, we use a variant of model (2) to compute an energy-efficient network in each demand scenario $t \in T$. Again we fix the physical topology and the installed IP routers by fixing the variables y_e and x_i^n to the values of the base network, and minimize the number of line cards. In contrast to the previous case, not only the flow variables can be changed, but also the capacity variables y_p . In order not to exceed the number of line cards installed at each node in the base

network, we add the constraints $\sum_{p \in P_i} y_p \leq \sum_{p \in P_i} y_p^{base}$ for every node $i \in V$.

IV. DATA

We have made an effort to use as realistic data (network topologies, traffic demands, costs, and power) as possible. We have used the detailed hardware and cost model for IP and WDM equipment from [15], which has been developed by equipment vendors and network operators within the European NOBEL project [16].

Cost and power of network elements Every network node can be equipped with one out of nine different IP routers accommodating 16–208 line cards with a capacity of 640–8320 Gbps. Routers with a capacity of more than 640 Gbps are multi-chassis configurations. We considered a 40 Gbps colored line card interface that connects the IP router to the WDM system. To estimate the cost of this interface following [15], we combined a 40 Gbps IP router slot card, a 4x10 GE LR port card, and a 4x10G ELH muxponder. The power was evaluated by combining a Cisco 4-port 10-GE Tunable WDM PHY PLIM and a Modular Services Card which together consume 500 W [2, 17].

We assume an 80-channel optical system. Following [15], an optical fiber installed on a physical link is composed of optical line amplifiers (OLA), dynamic gain equalizers (DGE), dispersion-compensating fibers (DCF), and WDM multiplexers. The corresponding total cost of the fiber depends on the length of the physical link. As in [15] we assume an OXC to be composed of wavelength-selective switches (WSS), which results in a fixed cost and a cost that linearly scales with number of connected fibers. We may hence map the latter to the cost of fibers.

Notice that out of the 80 available wavelength channels per fiber, at most 57 were used in the networks resulting from our computations, even for high demand values. Assigning wavelengths to the lightpaths in a postprocessing step should thus not require the installation of any wavelength converters.

Network topology We used a German backbone network (ger17) with 17 nodes and 26 links that has been defined as a reference network in the NOBEL project [16] (Fig. 2). For every node pair (i, j) we precalculated the set $P_{(i,j)}$ of the 50 shortest paths using the spherical distances between the nodes. The paths were limited in length to 1500 km.

Demands We have evaluated the possible energy savings with different temporal and spatial demand distributions, on different time scales (every 15 min, every day, every month), and with different maximum total demands (1, 3 and 5 Tbps). One set of demands for ger17 was taken from measurements in the national research backbone network operated by the German DFN-Verein [18]. These have been taken in 5-minute intervals in the year 2005. These matrices were aggregated to time intervals of 15 minutes over the day 2005-02-15, intervals of one day over February 2005, and intervals of one month over year 2005.

To vary the ratio between demands and the capacity granularity, we scaled all demands by the same factor such that

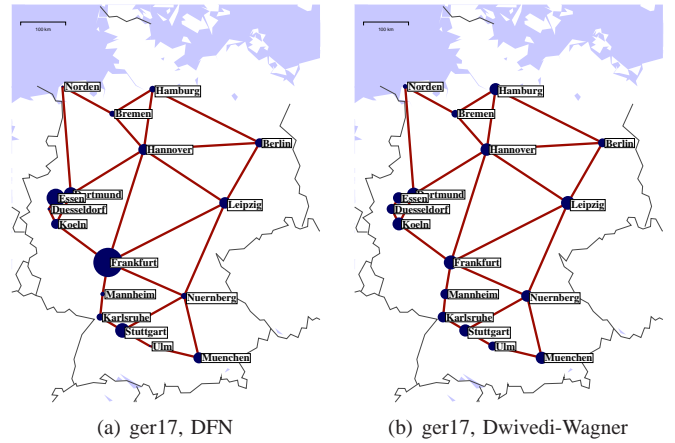


Figure 2. Physical topology, and source traffic distribution. The size of a node represents its emanating demand. The DFN demands (a) are Frankfurt-centralized in contrast to the DWG demands (b).

the sum $\sum_{i < j} d_{ij}^{max}$ of all demands in the maximum demand matrix was 1 Tbps, 3 Tbps or 5 Tbps. We refer to these values as the *maximum total demand*, while the value $\sum_{i < j} d_{ij}^t$ denotes the *total demand* at time $t \in T$.

The DFN matrices have a centralized structure with a large demand emanating from Frankfurt, which is a large entry point for cross-atlantic traffic. They also exhibit temporal peaks caused by single academic institutions sending large amounts of traffic to another institution or to an international backbone. Therefore we also evaluated the energy savings in ger17 with demand matrices generated using the Dwivedi-Wagner (DWG) model [19] based on population statistics. The resulting demands are much less centralized. Fig. 2 highlights the differences between the DFN and DWG demands; the area of each node is proportional to its emanating demand. From the single demand matrix $(b_{ij})_{(i,j) \in V \times V}$ obtained from the DWG model, we have generated demand matrices for all time periods by applying the relative demand changes in the DFN measurements to the computed DWG matrix as follows. Given the DFN demands (d_{ij}^t) over time, the maximum DFN demands (d_{ij}^{max}) , and the single DWG matrix (b_{ij}) , we calculate dynamic DWG demands b_{ij}^t in the following way:

$$b_{ij}^t := b_{ij} \cdot d_{ij}^t / d_{ij}^{max}.$$

The time-dependent scaling factor d_{ij}^t / d_{ij}^{max} takes values in the interval $[0, 1]$, which for every (i, j) normalizes the maximum DFN demand. It hence rules out the domination effects caused by single demands in the measurements. Fig. 3 illustrates this effect for the daily total demand values over one month.

V. RESULTS

The constructed base network is over-provisioned by definition since peak demands do not occur all at the same time. Consequently, allowing to change the IP-routing or the realization of lightpaths in low-demand hours with the objective of minimizing the number of line-cards has a dramatic effect

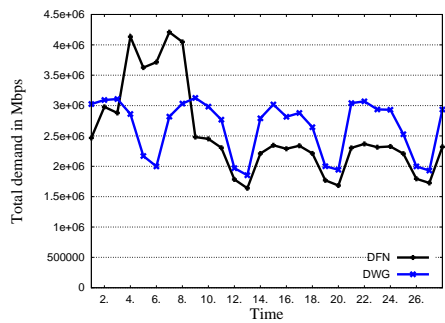


Figure 3. Total demand over time, every day of February 2005, maximum total demand value 5 Tbps. The DWG matrices show the expected behavior over time with peaks during the week and low traffic on weekends. In contrast, the DFN measurements exhibit peaks caused by single demands.

on the consumed energy. One should hence not overestimate the total energy savings compared to the static base scenario. Although these values are still interesting from a practical perspective, the more important comparison is the one between the three strategies FUFL, DUFL, and DUDL in low- and high-demand situations.

All computations have been carried out on an 64-bit Intel 3.00 GHz CPU with 8 GB main memory. All occurring MILPs have been solved using CPLEX 12.1 [20] as a black-box solver with a time limit of 1 hour.

In the following we focus on the results for the ger17 network and the 96 DFN traffic matrices given for every 15 minutes of 24 hours, with 1, 3 and 5 Tbps maximum total demand. The observed phenomena were consistent over all considered scenarios (28 DFN matrices over a month, 12 DFN matrices over a year, DWG demands).

For most of the DUFL optimizations an energy-minimal solution was obtained within seconds or minutes. For only a few instances we hit the time-limit with an optimality gap (relative difference between the number of line cards in the best solution and a mathematically proven lower bound to this number) below 5%. The optimization problem corresponding to DUDL is harder to solve. All DUDL runs hit the time-limit with optimality gaps of 11%–30% (1 Tbps) and 3%–15% (5 Tbps). All comparisons of the three strategies will be made against the lower bound on the number of line cards in use, which corresponds to an upper bound on the maximum possible energy savings in the considered scenario. Note that for FUFL and for almost all DUFL runs dual bounds and solution values are identical.

Figures 4(b)–(d) illustrate the solution values for the base network (constant over time), the power values obtained with FUFL, and the values and dual bounds for DUFL and DUDL. Notice that these values relate to power of line cards only, which constitute 23.1/22.5/24.5 % of the total power of line cards and IP routers for the static base network 1/3/5 Tbps of maximum total demand, respectively.

The power consumption generally follows the total demand curve (Fig. 4(a)). In the 1 Tbps scenario, the network consumes 0.77/0.57/0.41 MWh over the day for FUFL/DUFL/DUDL,

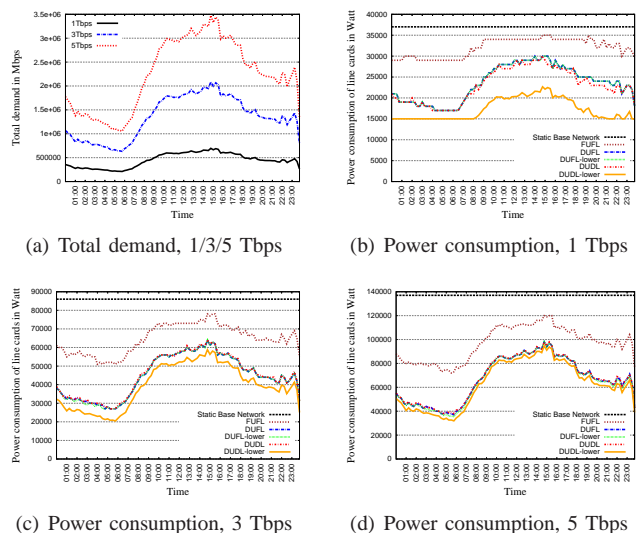


Figure 4. ger17 network, DFN traffic, 1/3/5 Tbps, every 15 minutes on February 15, 2005. The power consumption curves follow the total demand curve. The difference between FUFL and DUFL is much larger than the additional benefit of DUDL.

respectively. The corresponding values for 3 Tbps and 5 Tbps are 1.55/1.07/0.95 and 2.32/1.59/1.51 MWh. These figures show that reconfiguration in the virtual layer by DUFL saves much energy compared to FUFL. In contrast, reconfiguring the virtual topology in the physical layer (DUDL) does not give much additional profit. In the 5 Tbps scenario the DUFL and DUDL curves nearly coincide. The lower bound for DUDL proves that only a small amount of energy can be saved compared to DUFL. There seems to be more tolerance in the 1 Tbps scenario. In this case we cannot verify whether our DUDL solutions are optimal or whether there exist solutions closer to the lower bound.

Although comparing FUFL to the over-dimensioned base network (3.288 MWh consumed over the day by the line cards only) may lead to misinterpretations, we note that FUFL alone might give substantial savings and is easy to realize in practice. In the 5 Tbps scenario FUFL reduces the power of the active line cards in low-demand hours by up to 47 % at 05:30 am (74 % for DUFL and 76 % for DUDL). Even in the high-load scenario the savings are significant (12 %, 28 %, and 31 % for FUFL, DUFL and DUDL at 02:45 pm, respectively). Considering the power consumption at 05:30 am and 02:45 pm for a maximum total demand of 5 Tbps, 40 % of power for FUFL, 63 % for DUFL, and 66 % for DUDL can be saved in the early morning compared to the peak hour.

To understand the relatively poor outcome of DUDL, one has to consider two extreme scenarios. If the demand in the network is very large, the virtual topology in the base network is close to a full mesh. Since DUFL may use any virtual link from the base network, the DUFL solution is (close to) optimal. DUDL cannot benefit from choosing lightpaths inactive in the base network. If, on the other hand, the demands are very small, the optimal virtual topology of the base network will be a tree. Also both DUFL and DUDL will find a tree network.

Both trees might differ, but use the same number of line cards. Hence again DUDL cannot benefit compared to DUFL. For the ger17 network, Fig. 4(b) illustrates that DUDL saves energy only in peak demand hours of the 1 Tbps scenario.

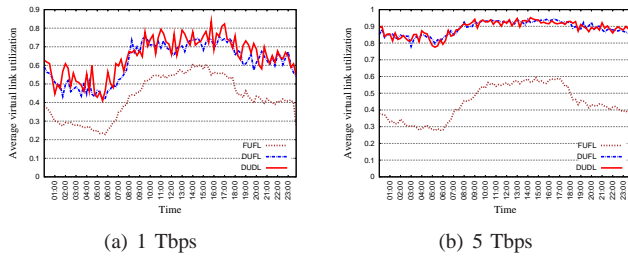


Figure 5. Average virtual link utilization over all active virtual links for ger17, DFN, 1/5 Tbps, every 15 minutes. DUDL and DUFL achieve high lightpath utilization due to dynamic multi-path IP routing.

For the success of DUFL it may be crucial that we admit flow splitting in the virtual domain, which allows DUFL to fill up the established lightpaths to a high extent (see Fig. 5(a) and (b)).

VI. CONCLUSIONS

Our study has shown that a significant amount of energy can be saved by switching off line cards in low-demand hours with any of the considered reconfiguration strategies FUFL, DUFL and DUDL. The formulated MILP allowed us to provide high quality solutions together with upper bounds on the maximum possible energy savings in the corresponding multi-layer optimization problems. We used realistic topologies, traffic data, cost and power values. Our main result is that rerouting demands in the IP layer (DUFL) contributes the most to the energy savings. Allowing additional reconfiguration in the optical domain (DUDL) barely brings any extra benefit in the considered scenarios. Reconfiguring the IP routing is nowadays part of the daily business of network operators. Our work indicates that energy should be included in the objective function whenever IP routing weights are reconfigured in the operational phase. It should also motivate equipment vendors to provide line cards with a convenient and fast functionality to be switched on and off.

ACKNOWLEDGMENT

This research has been supported by the DFG research Center MATHEON, and by the European Commission within the FP7 4WARD and BONE projects.

REFERENCES

- [1] United States Department of Energy, “International energy outlook 2008,” [http://www.eia.doe.gov/oiaf/ieo/pdf/0484\(2008\).pdf](http://www.eia.doe.gov/oiaf/ieo/pdf/0484(2008).pdf), 2008.
- [2] “Cisco CRS-1 Production Brochure,” http://www.cisco.com/en/US/prod/collateral/routers/ps5763/prod_brochure0900aecd800f8118.pdf, 2008.
- [3] 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*, 2008, pp. 1130–1138.
- [4] M. Gupta and S. Singh, “Greening of the internet,” in *Proc. of the ACM SIGCOMM, Karlsruhe, Germany*, 2003.
- [5] L. Chiaraviglio, M. Mellia, and F. Neri, “Reducing power consumption in backbone networks,” in *Proc. of the ICC, Dresden, Germany*, 2009.
- [6] —, “Energy-aware backbone networks: a case study,” in *Workshop on Green Communications, Dresden, Germany*, 2009.
- [7] Y. Wu, L. Chiaraviglio, M. Mellia, and F. Neri, “Power-aware routing and wavelength assignment in optical networks,” in *Proc. of the ECOC, Vienna, Austria*, 2009.
- [8] B. Puype, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, “Power reduction techniques in multilayer traffic engineering,” in *Proc. Of the ICTON, Azores, Portugal*, 2009.
- [9] B. Puype, D. Colle, M. Pickavet, and P. Demeester, “Energy efficient multilayer traffic engineering,” in *Proc. of the ECOC, Vienna, Austria*, 2009.
- [10] G. Shen and R. S. Tucker, “Energy-minimized design for IP over WDM networks,” *Journal of Optical Communications and Networking*, vol. 1, no. 1, pp. 176–186, 2009.
- [11] F. Ortega and L. A. Wolsey, “A branch-and-cut algorithm for the single-commodity, uncapacitated, fixed-charge network flow problem,” *Networks*, vol. 41, pp. 143–158, 2003.
- [12] S. Terblanche, R. Wessäly, and J. Hattingh, “Solution strategies for the multi-hour network design problem,” in *Proc. of the INOC, Spa, Belgium*, 2007.
- [13] A. Bley, U. Menne, R. Klähne, C. Raack, and R. Wessäly, “Multi-layer network design – a model-based optimization approach,” in *Proc. of the PGTS, Berlin, Germany*, 2008, pp. 107–116.
- [14] S. Raghavan and D. Stanojević, “WDM optical design using branch-and-price,” 2007, working paper, Robert H. Smith School of Business, University of Maryland.
- [15] R. Hülsermann, 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.
- [16] “NOBEL – Next generation Optical networks for Broadband European Leadership,” <http://www.ist-nobel.org>, 2004–2007.
- [17] F. Idzikowski, “Power consumption of network elements in IP over WDM networks,” TU Berlin, TKN Group, Tech. Rep. TKN-09-006, 2009.
- [18] “Deutsche Forschungsnetz Verein,” <http://www.dfn.de>, DFN.
- [19] A. Dwivedi and R. E. Wagner, “Traffic model for USA long distance optimal network,” in *Proc. of the OFC, Baltimore, USA*, 2000, pp. 156–158.
- [20] IBM – ILOG, “CPLEX,” <http://www.ilog.com/products/cplex/>, 2009.