Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

Filip Idzikowski[1]
Sebastian Orlowski[2]
Christian Raack
Hagen Woesner[1]
Adam Wolisz[1]

Dynamic routing at different layers in IP-over-WDM networks – Maximizing energy savings

[1] TKN, TU Berlin, Germany
[2] atesio GmbH, Berlin, Germany

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

ZIB-Report 10-07 (May 2010)

Dynamic routing at different layers in IP-over-WDM networks – Maximizing energy savings *†

Filip Idzikowski
TKN, TU Berlin, Germany

Sebastian Orlowski
atesio GmbH, Berlin, Germany

Christian Raack
ZIB, Berlin, Germany

Hagen Woesner
TKN, TU Berlin, Germany

Adam Wolisz
TKN, TU Berlin, Germany

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 three realistic network topologies 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. These results are independent of the ratio between the demand and capacity granularities, the time scale and the network topology, and show little dependency on the demand structure.

Keywords: network design, energy efficiency, power consumption, multi-layer, multi-hour, multi-period

1 Introduction

In the light of scarce resources and the rising demand for energy there is a growing interest in solutions and "green" strategies in different fields to reduce the power consumption[42]. In this work, we focus on energy efficiency in IP (Internet Protocol) over WDM (Wavelength Division Multiplexing) backbone networks.

Telecommunication networks are typically dimensioned to handle an estimated worst-case traffic scenario. The classical approach to network design hence assumes a given single traffic matrix, see [7, 8, 23, 26, 32] and the references therein. This matrix can be estimated, e.g., by using population statistics [13] or by exploiting information from current traffic measurements [8, 49]. To handle future changes in the traffic volume and potential peaks in high demand hours, the demand values are very often (highly) overestimated. This conservative approach leads to static solutions supporting all potential traffic patterns but results in overprovisioned networks and a waste of CAPEX (capital expenditures) as well as OPEX (operational expenditures). There have been several attempts in the literature to handle traffic uncertainty already in the design process stemming from stochastic or robust optimization [4, 25, 27, 31].

This work rather focuses on traffic engineering, given a statically designed, capacitated multi-layer network. Our aim is to dynamically adapt the routing and the number of active components to the traffic patterns reducing OPEX, where energy is one of the key factors. It is a well-known fact that the power

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

†A short version of this paper has been published in the proceedings of the ONDM 2010 [20]

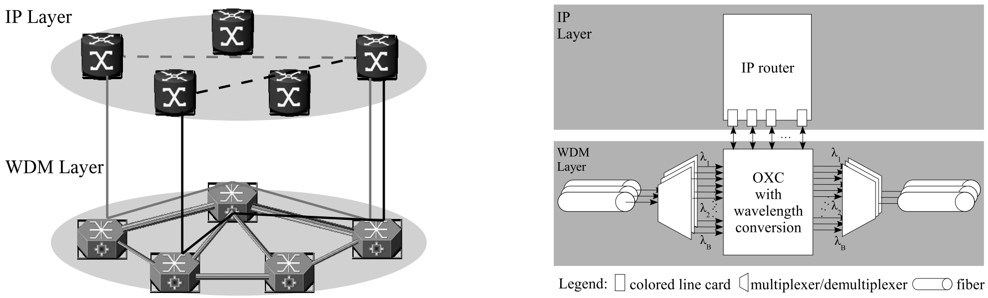

(a) Network model

(b) Node architecture

Figure 1: IP-over-WDM network: optical cross-connects interconnected by fibers constitute the physical topology, and IP routers interconnected by lightpaths constitute the virtual topology

consumption of IP routers and line cards is almost independent of the load and may reach thousands of kW [1, 9, 16]. We pose the question of how much energy can be saved by dynamically switching off idle IP router line cards in low-demand hours. In this paper we 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. These approaches correspond to three different levels of freedom to dynamically change the routing in the WDM and IP layers. They are presented in Section 2 together with the used network model.

Although several papers have focused on power consumption in single- and multi-layer networks (see Section 3), 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 4, 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 provable energy-optimal solutions or at least upper bounds on the potential savings. Section 5 describes the used data. The computational study in Section 6 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 7 concludes our work.

2 Network model

We focus on IP-over-WDM networks, where the WDM layer offers optical by-pass technology, as depicted in Fig. 1 (a). Nodes in the WDM layer, which represent optical cross-connects (OXCs), 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 (see Fig. 1 Fig. 1(b)), which provide a direct interface between IP and WDM by performing 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 (of capacity corresponding to the number of lightpaths between these routers) in the IP layer. The virtual links together with the IP routers form a virtual topology. All IP routers are sources and destinations of aggregated backbone traffic, which is converted into an optical signal by the line cards and directly fed into OXCs. We assume full wavelength conversion capability, and leave the wavelength assignment and installation of converters to a postprocessing step [21].

A lightpath between two particular line cards may be realized (routed) by different physical paths in the WDM layer. Similarily, the IP traffic demands can be routed over different virtual links. 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 using different

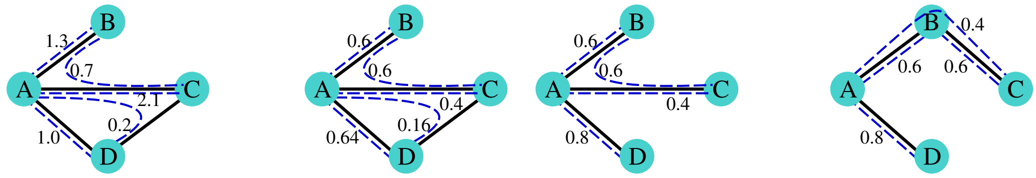

(a) base network, 7 lightpaths, 14 line cards (b) FUFL, 6 lightpaths, 12 line cards (c) DUFL, 4 lightpaths, 8 line cards (d) DUDL, 3 lightpaths, 6 line cards

Figure 2: 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.

physical paths).

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 changing 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 in the fiber network.

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 at routing 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 installed line 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 changes of the virtual topology in this case.

Fig. 2 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 not touched. 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 [15, 29]. DUFL is a single-layer network design problem which can be solved to optimality in a reasonable amount of time in practice, see Section 6 and [7, 26]. DUDL is a computational challenge since it involves optimizing two coupled network layers simultaneously, similar to designing the base network, see [12, 14, 23, 24, 28].

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 involved routers. Even more signaling is needed for additional 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. Generalized Multiprotocol Label Switching (GMPLS) is a potential candidate for managing the paths. The reconfiguration itself is non-trivial though since it requires the use of OXCs to dynamically change virtual links (typically realized by point-to-point connections nowadays). The energy-efficient network design may have an influence on resilience and QoS (Quality of Service) in the network in terms of packet delay, jitter, packet loss as well as on network throughput, since switching off line cards decreases capacity of the network. Moreover, network devices being repeatedly switched on and off may be more prone to failures. 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.

3 Related work

Although energy saving is quite a new subject in the wireline networking research, it has already been addressed in quite a few papers since the pioneering work by Gupta and Singh [16]. We focus on routing to save energy and divide the papers into three groups depending on the dynamics of routing at different network layers. In this literature review, we look at the source of the saving potential (what kind of network elements can be switched off), the approach taken to determine the energy savings (analytical, optimization or simulation), the considered scenarios (topology, traffic, power and cost values) and the dynamics of the network over time.

Dynamic routing in the optical layer The Power-Aware Routing and Wavelength Assignment (RWA) problem is proposed by Wu et al. [46] and formulated as a MILP. Energy savings can be achieved by switching off OXCs and optical amplifiers according to three proposed algorithms. A rough lower bound is also presented. Investigations of bidirectional rings and generic meshes of up to 32 nodes without wavelength conversion, and with a large number of fibers on each link and of wavelengths per fiber compared to the number of lightpaths (random requests with no dependency on time) revealed that smart routing of lightpaths in the WDM layer may bring significant energy savings against the shortest path routing and first fit wavelength assignment. The authors assume that both the power of an amplifier and the power of an OXC are equal to 1 kW (unreferenced).

Silvestri et al. [39] make use of traffic grooming and transmission optimization to reduce energy consumption in the WDM layer. Traffic grooming shifts traffic from some links to other ones in order to switch empty ones off, and transmission optimization adjusts dispersion management and pulse duration which decreases the need for using in-line 3R regenerators. Taking the simulative approach (OPNET SP Guru Transport Planner), the authors consider a European transport network (26 nodes and 46 links) with and without OXCs, and scale a traffic matrix to mimic changing demands. The power consumption of an optical amplifier is estimated to 200-500 W and the power of a 3R regenerator to 2-5 kW (unreferenced). The results show that transmission optimization may lead to the elimination of in-line 3R regenerators, and that traffic grooming allows to switch off significant number of links.

Bathula et al. [5] leverage node clustering and anycast routing to obtain a trade-off between the energy consumption and the average requests lost due to the sleep cycles of nodes in the clusters. In order to allow clusters to be switched to an OFF state, requests are destined to a set of nodes. If a destination cannot be reached due to its intermediate node belonging to a cluster in an OFF state, the next available destination can be chosen under acceptable bit-error-rate (BER) and propagation delay. The proposed algorithm is applied to the NSF network (14 nodes and 21 links) with four clusters, where call arrivals follow a Poisson process. Total power is calculated using the energy per bit for a core wavelength routing node (WRN) and an optical amplifier such as EDFA (approximately 10 nJ and 0.1 nJ respectively) as well as the power of a transmitter and a receiver (unreferenced). The authors discuss the simulation results in the light of a trade-off between the average power consumed for each request and average request blocking as functions of load in Erlangs and number of clusters in an OFF state.

Dynamic routing in the electrical layer Electrical network components offer a high potential of energy savings due to their high energy consumption [19, 30, 37]. Usage of dynamic routing in the electrical layer for energy saving was investigated in several papers. Many of the approaches therein share their concept with DUFL.

Chabarek et al. [9] introduce power aware network design, where power consumption is reduced by adjusting routing of flows, choosing appropriate chassis type at each node in terms of capacity and power, and allocating appropriate number and type of line cards at each chassis. The proposed MILP minimizes power of all the chassis and line cards in the network. The authors support their approach with measurements of power consumption of routers (Cisco GSR 12008 and Cisco 7507). They point out that the power consumed by the routers shows only little dependency on the load (data rate, packet size, packet inter-arrival times). Despite router chassis being the most power hungry elements of the router, its power consumption is highly dependent on the number and type of installed line cards. Using the power measurements and the MILP the authors investigate power consumption of 3 random and 4 Rocketfuel networks (7-21 nodes, 18-134 links) with traffic matrices generated according to the gravity model. The traffic is scaled with several factors in order to observe changes of power consumption in dependence of the load. The dynamics of traffic over time is not considered. After relaxing some constraints, the authors find out that minimum power consumption coincides with chassis that can accommodate a large number of line cards and line card capacities that closely match demand.

Chiaraviglio et al. [10, 11] focus on power consumption in hierarchical networks, where it is possible to turn off nodes and links. They consider a single layer routing problem with time-varying demands, model it with an MILP [11], and use numerous heuristics [10, 11] to solve it. A randomly generated network topologies consisting of 10 core nodes, 30 edge nodes and 120 aggregation nodes are considered in [11], and a network similar to the one of the largest ISPs in Italy consisting of 8 core nodes, 52 backbone nodes, 52 metro nodes and 260 feeders is considered in [10]. A traffic matrix between the aggregation nodes [11] an feeders [10] is randomly generated. In order to mimic the dynamics of traffic over time, the traffic matrix is scaled by a sinusoidal [10, 11] and the ISP's profile [10]. Time intervals of 5 minutes over a day are considered. The authors show that significant power can be saved. While [11] reports on the percentages of nodes and links that can be switched off, [10] assumes some power values of amplifiers and router interfaces (unreferenced).

Energy Profile Aware Routing (EPAR) introduced by Restrepo et al. [36] assumes the dependency between the energy consumption and the traffic load or traffic throughput of a particular network component, which is referred to as the Energy Profile. Energy can be saved by choosing components with appropriate Energy Profiles. EPAR (formulated as a linear equation system) using five Energy Profiles for routers (energy consumption range 0 - 15 kWh under 0 - 3.2 Tbps traffic, respectively) is evaluated on a Germany50 network (50 nodes, 88 links) under fully-meshed traffic (no further details about the traffic available, traffic dynamics not considered). Shortest (least hop) path routing is used as a reference. Significant energy savings are shown. Cubic Energy Profiles turn out to be especially efficient. Convex profiles achieve higher energy savings (against the shortest path routing) than the concave ones due to multi-path routing making use of low power consumption over broader load range.

Energy-Aware Traffic Engineering (EATe) introduced by Vasić and Kostić [43, 44] is another approach for energy saving. Changing routes aims not only at switching off links and routers, but also at rate adaptation. The authors evaluate EATe on five Rocketfuel topologies (19-115 nodes, 68-296 links) using ns-2 simulations [43], TRUMP traffic and given drop margin, which determine the number of links that EATe tries to push to a lower energy level. Four uniformly distributed energy operating rates with quadratic energy savings between them are used and relative energy savings are reported. It is shown that EATe manages to completely remove traffic on up to 31% of the links (depending on the number of alternative paths for the traffic in the network), without a significant increase of link utilizations. It is also possible to put routers to sleep with little increase of link utilizations, as well as make use of rate adaptation. Based on link utilizations the authors predict that most of the time EATe will have negligible impact on latency. Stability of EATe under traffic changes and handling of link failures are also presented.

Puype et al. [33, 34] investigate multi-layer traffic engineering with the objective of reducing the power consumption in IP over optical networks. They assign higher routing costs to low-load logical links in order to empty them after PI rerouting. Empty logical links are removed from the actual logical topology to save energy. Since routing of logical links is not considered, we classify this work as dynamic routing in the electrical layer. Applying an algorithmic approach to a 14-node network under a traffic pattern based on a uniform distribution, the authors show significant power savings against a full-mesh logical topology.

Traffic characteristics cover diurnal traffic variations - off-hour traffic is equal to 0.25 of the peak traffic). Authors investigate the case when logical topology updates are slower than the diurnal traffic variations and the opposite one under the assumption that power of router interfaces (line cards) depends on the carried traffic. The authors discuss two ways to influence the power vs. bandwidth (load) curve, i.e. idle power reduction (by e.g. matching line rates with traffic volume, or reducing clock rates) and the scaling of equipment power requirements using newer CMOS technology. Both of them influence the results of the considered traffic engineering approach. Normalized power values against the maximum power are reported.

Dynamic multi-layer routing The possibility to change the routing in both the IP layer and WDM layer intuitively offers the greatest opportunities for energy savings.

Extensive work on power efficient networks has been done by the group of Prof. Tucker. [37] is especially relevant to our work. It tackles IP-over-WDM backbone transport networks with no OXC switching capability. The authors formulate a MILP to minimize power consumption of the network in two layers (router ports, transponders, and Erbium-doped Fiber Amplifiers can be switched off) and propose two heuristics for energy efficient IP-over-WDM networks (”direct bypass” and ”multi-hop bypass”). They consider three networks: a test network with 6 nodes and 8 links, the NSFNET network with 15 nodes and 21 links, and USNET with 24 nodes and 43 links. The heuristics are compared against the optimum solution (for two smaller networks), ”non-bypass” network and LP relaxed MILP (for all the networks). In the ”non-bypass” no IP router can be bypassed by a lightpath. The relaxation of a MILP is achieved by allowing all the integer variables to take real values. Assuming static uniformly distributed random traffic and realistic power values for the considered networks, the authors conclude that the IP routers contribute the most to the power consumption of the network. Moreover, they point out that the power savings achieved by the lightpath bypass strategy increase with the network size, and that the power consumption distribution across the network nodes has a large variance for the ”non-bypass” case, as opposed to the ”direct bypass” and ”multi-hop bypass”. Eventually the authors note that minimizing energy and minimizing cost gives similar results in all the considered scenarios.

Yetginer and Rouskas [47] consider a two-layer architecture with each node equipped with an OXC and a DXC. Routing of lightpaths over the physical layer (fibers) and routing of traffic over the lightpaths are taken as variables in the proposed ILP. Following the metric proposed in [9], the authors of [47] define the power consumption of the network as a weighted sum of the number of lightpaths and total amount of traffic electronically routed. Three objective functions originate out of this metric: Minimum Number of Active Router Ports (minL), Minimum Amount of Electronically Switched Traffic (minT) and Minimum Power Consumption (minP). A 6-node network with 8 links under random traffic (uniform distribution with varied maximum value) is studied. The power consumption of a lightpath under no load (fixed power consumption) is equal to 0.25 of the power consumption of a lightpath under full load (maximum power consumption). The results indicate that minP uses only a few more lightpaths than minL. The difference between minP and minT in terms of electronically switched traffic vanishes as the network load is increased. Traffic dynamics over time is not discussed.

A trade-off between energy efficiency and CAPEX minimization is studied by Palkopoulou et al. [30]. Two network architectures are considered: IP-over-WDM (IPoWDM) and IP-over-optical-transport-network-over-WDM (IPoOTNoWDM). Energy efficiency is optimized (the optimization model is not presented in the paper) in the Germany17 network under random traffic (uniform distribution with varied maximum value, but no dynamics over time). Transport link failures and core router failures are taken into account. Cost values taken from [17] and power values according to internal Nokia Siemens Networks estimations apply. Transponders, router port cards and EXC port cards contribute to the power consumption with the dominating contribution of the IP equipment. The investigations show that although minimization of CAPEX coincides with minimization of energy in terms of power consumption with changing load and fixed network architecture, the most cost efficient architecture is not always the least energy consuming one under certain load. The optimal architecture in terms of power consumption is dependent on the inter-node traffic demand. However, the relative power contribution of different network layers is independent of the average inter-node traffic demand for both IPoWDM and IPoOTNoWDM. A prediction on future power of network equipment is also made.

Routing in IP and optical layers is also considered by Shen et al. [38]. The ILP proposed by the authors has two objective functions, minimizing power or cost of the network. The same 6-node 8-link network as in [30] is considered. It is fed with random traffic (uniform distribution with varied maximum value). The same values for power [9] and cost are applied, with the only difference that each processed traffic unit consumes additional power, but contributes no costs though. The authors show that multi-layer networks consume approximately 80% of the total power consumption of networks with no bypassing of routers (referred to as IP networks). The authors also show that the profit of diversified-volume lightpaths in terms of network power-efficiency against nondiversified-volume lightpaths decreases with increasing load. Eventually, the power-minimized network is shown to be more power-efficient than the cost-minimized network due to penalty for traffic processing. Dynamics of traffic over time is not taken into account.

Our contribution Various approaches to energy saving on different layers have been proposed so far in the literature. They all show that significant energy savings can be achieved, however it has been stressed in many papers that the presented MILPs are NP-hard, and therefore small networks or heuristics are considered. Moreover, random traffic data and unreferenced power and cost values have been used in some papers. Our work extends the work presented above in that we consider and compare several rerouting strategies on different layers in the IP-over-WDM network (FUFL, DUFL, DUDL). We evaluate the potential energy savings using sophisticated mixed-integer programming methods with time-varying traffic demands obtained from measurements in realistic networks, using realistic cost and power consumption values of the network equipment. We cover different time-scales, temporal- and spatial distribution of traffic, and demand scalings. We find solutions, which are mostly optimal. Moreover, we have found no previous work, which tackles FUFL ([39] is the closest in the WDM layer).

4 Methodology and mathematical models

In the first step we design an IP-over-WDM network including installation of fibers and hardware. This network serves as a basis for our investigations, and is considered to be *static*: It is 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 possible peak demands in a given period of time simultaneously without any reconfiguration. Among all possible topologies and routings with this property, 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 minimizing the number of active line cards.

4.1 Design of a base network

Designing a cost-minimal multi-layer network that allows to realize a given demand matrix is a highly complex problem which is far from being solved yet, see for instance [12, 14, 23, 24, 28]. The problem becomes even harder if multiple demand matrices are to be considered, leading to robust multi-layer network design problems. To the best of our knowledge the only study considering multi-layer networks and robustness can be found in [6]. Most attempts to solve robust network dimensioning problems, however, assume single-layer networks, see [4, 27, 40] and references therein.

Given detailed traffic measurements, our approach is based on constructing a single demand matrix that refers to all peak demands over time. The base network is then dimensioned with respect to this maximum matrix which ensures that every single traffic scenario can be realized. The base network can be considered being cost-minimal among all networks that allow to route the constructed maximum matrix. Note that although common in practice, this approach is potentially producing overprovisioned networks. There might be cheaper topologies that cannot accomodate the maximum traffic matrix but all single traffic scenarios (with a static IP routing). It is also typically unlikely to have all demands at their peak simultaneously. For our study, however, this approach based on a maximized matrix is reasonable since our main goal is to evaluate different energy saving principles rather than providing the cheapest among all robust base networks.

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

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

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

Our model used to cost-optimally design the base network is close to the one used in [8, 22, 23, 28, 35]. We optimize both network layers at the same time in an integrated step. The model 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 $i \in V$ 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 simplicity and due to the absence of realistic power data we do not consider different optical nodes to be available. Instead we assume a preinstalled optical cross-connect (OXC) of infinite capacity at every network node. The model (2) can be easily extended by a physical node model though. Also note that we include some cost for OXCs in the cost for fibers, see Section 5.

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. Every path $p \in P$ can be equipped with multiples of the bandwidth C. Each of these bandwidth units C on a path p incurs the cost γ of two line cards at the end-nodes of p and 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 [7] reduces the number of commodities from $\mathcal{O}(|V|^2)$ to $\mathcal{O}(|V|)$ and leads to commodities with 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, i.e. the demand value emanating from the common source of its demands, 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 the virtual link between i and j in both directions. These variables aggregate the IP traffic on all lightpaths with end-nodes i and j. Notice that these flow variables are defined between pairs of nodes at the virtual layer and not on individual lightpaths. This is possible because the actual physical representation of a virtual link does not matter for the IP routing. Only the total capacity between any two nodes is of interest for an IP demand in our model. 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, as explained below and in [28, 35]. This significantly reduces the size of the model compared to considering flow variables on individual physical representations. 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, formulated on the complete virtual layer $V \times V$. 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.

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

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

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

$$\sum_{n\in N} C^n x_i^n - \sum_{p\in P_i} Cy_p \ \geq \ d_i, \ \ i \in V \tag{2c}$$

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

$$By_e - \sum_{p\in P\colon e\in p} y_p \ \geq \ 0, \quad e \in E \tag{2e}$$

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

The flow for commodities $k \in K$ on virtual link $(i,j) \in V \times V$ has to be disaggregated in two steps in postprocessing. First, we have to disaggregate the flow from the virtual links (i,j) to all physical representations $p \in P$ with end-nodes i and j (in both directions). By (2b) this can be done respecting the installed capacities Cy_p for $p \in P_{(i,j)}$. Let f_p^k denote the flow on path p for commodity $k \in K$. In a second step, the commodity flow f_p^k has to be disaggregated to a flow for every individual point-to-point demand $(i,j) \in V \times V$, where i corresponds to the source of commodity k and j to one of its targets. Since commodities have been aggregated from demands at common source nodes, this can be done in a greedy manner, see e.g. [7]. We denote by $f_p^{(i,j)}$ the flow for demand (i,j) with value $d_{ij}^{\max}$ on path p.

4.2 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 end-to-end path for this demand by the factor $d_{ij}^t / d_{ij}^{\max} \in [0,1]$ and reduce the capacity on the path accordingly. As all flows for this demand are scaled by the same factor, the relative share of traffic on each used path remains the same as before.

To state it more precisely, consider a path p used to route demand $d_{ij}^{\max}$. In scenario $t \in T$, we reduce the flow from $f_p^{(i,j)}$ to

$$f_p^{(i,j)^*} := f_p^{(i,j)} \cdot d_{ij}^t / d_{ij}^{\max}.$$

Let

$$f_p{}^* := \sum_{(i,j)\in V\times V} f_p^{(i,j)^*}$$

be the total flow on path p after reducing the flow for all demands. Then the capacity on that path can be reduced from y_p to

$$y_p{}^* := \left\lceil \frac{f_p{}^*}{C} \right\rceil.$$

Remember that variables y_p count the number of lightpath on $p \in P$. Hence this procedure corresponds to shifting traffic between parallel lightpaths, that is, lightpaths with the same realization $p \in P$. No optimization is needed. From a practical perspective an alternative FUFL principle might be of interest. Traffic could be shifted between all paths $p \in P$ having the same end-nodes (i,j) ignoring physical representations. This could reduce the number of active lightpaths even more and completely relies on information available at the IP-layer. For the ease of exposition we neglect to study this FUFL variant here.

DUFL For every $t \in T$, we compute a new IP routing by using a variant of the model (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 compared to the base network, but not augmented. This is ensured by adding the constraint $y_p \leq y_p^{base}$ 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 without any further restrictions 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 the number of active lightpaths $\sum_{p \in P} y_p$.

DUDL With DUDL, we are allowed to change the IP routing as well as the virtual topology including physical representation of the lightpaths. But we cannot install new line cards at IP routers. 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 variables y_e (physical network) and x_i^n (IP routers) to the values of the base network, and minimize the number of line cards. In contrast to the previous case, augmentation of virtual link capacity variables y_p is allowed in addition to changing the flow variables. 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 $i \in V$. We then minimize the number of used line cards for every demand scenario using the same objective as for DUFL.

5 Data

We have made an effort to use as realistic data (network topologies, traffic demands, costs, and power) as possible. We have performed an extensive search for power values in the literature and product data sheets [19]. We have used the detailed hardware and cost model for IP and WDM equipment from [17], which has been developed by equipment vendors and network operators within the European NOBEL project [3]. Traffic data originates from measurements and determines network topologies, as reported further on in this section.

Cost and power of network elements In the following we briefly describe the network elements we used to design the IP-over-WDM architecture (see Fig. 1(b)). 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 [17], 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 WDMPHY PLIM and a Modular Services Card which together consume 500 W [1, 19].

We assume an 80-channel optical system. Following [17], 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 [17] 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.

Network topologies We used three network topologies as depicted in Fig. 3. The German backbone network Germany17 with 17 nodes and 26 links (Fig. 3 (a) and (b)) has been defined as a reference network in the NOBEL project [3]. The Abilene network (12 nodes and 15 links, Fig. 3 (d)) is an American network commonly used in the research community [48]. The biggest network we investigated (22 nodes and 36 links) is the pan-European research network Géant, which connects European National Research and Education Networks (NREN) [41]. We treated two German nodes as one due to location problems in contrast to the original data [41].

Considering each single network we precalculated the set $P_{(i,j)}$ of the 50 shortest paths for potential lightpaths for every node pair (i, j). The length of a physical link was computed by using the spherical

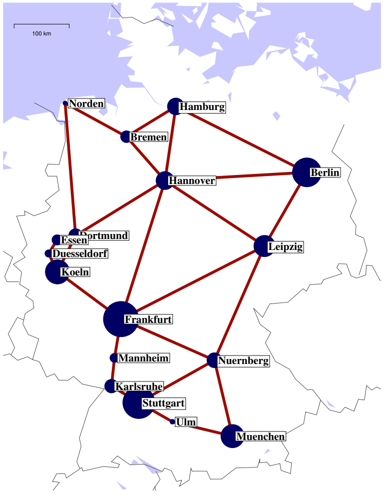

(a) Germany17, DFN

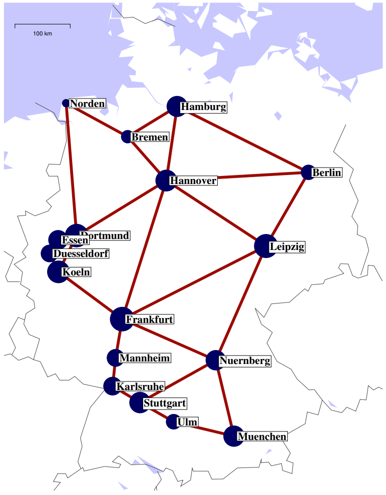

(b) Germany17, Dwivedi-Wagner

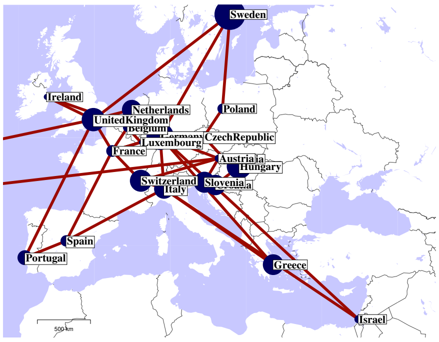

(c) Géant

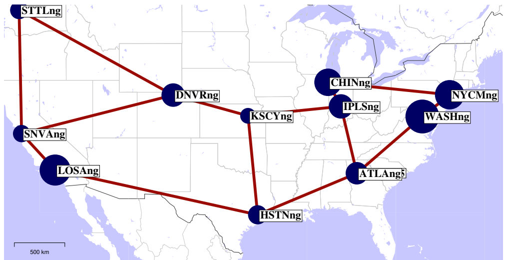

(d) Abilene

Figure 3: Potential 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). Notice that the New York node is omitted in the Géant figure (d).

distance of its end-nodes. The paths were limited in total length (sum of the physical link length) to 3000 km. Longer paths are rejected. There are three physical links in the Géant network with a real length greater than 3000km (Israel to Netherlands, New York to Austria and New York to the UK). We set these links to a length of exactly 3000km. There is thus only the direct path possible between these nodes. This way we implicitely assume to provide the necessary regenerators on these three links at no cost.

Traffic demands The choice of the network topologies presented above was dictated by the limited availability of the traffic measurements data. To vary the ratio between demands and the capacity granularity, we scaled all demands by the same factor such that 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 $t \in T$.

One set of dynamic demands (for Germany17) was taken from measurements in the year 2005 in the national research backbone network operated by the German DFN-Verein [2], see Fig. 4 (a), Fig. 5 (a), and Fig. 5 (b). The original DFN data consists of the total end-to-end traffic in Byte every 5 minutes over the day 2005-02-15, every day of February 2005, and every month over the year 2005. According to our partners at DFN-Verein, the traffic patterns in these periods were rather representative. We aggregated the 5-minute traffic matrices to 15-minute traffic matrices by taking the maximum value for each demand over the whole aggregation interval (in contrast to [20]). Eventually all matrices were mapped from the original DFN locations to the Germany17 network according to the smallest geographical distances.

A different set of traffic matrices was used for the Abilene network. It originates from the Abilene

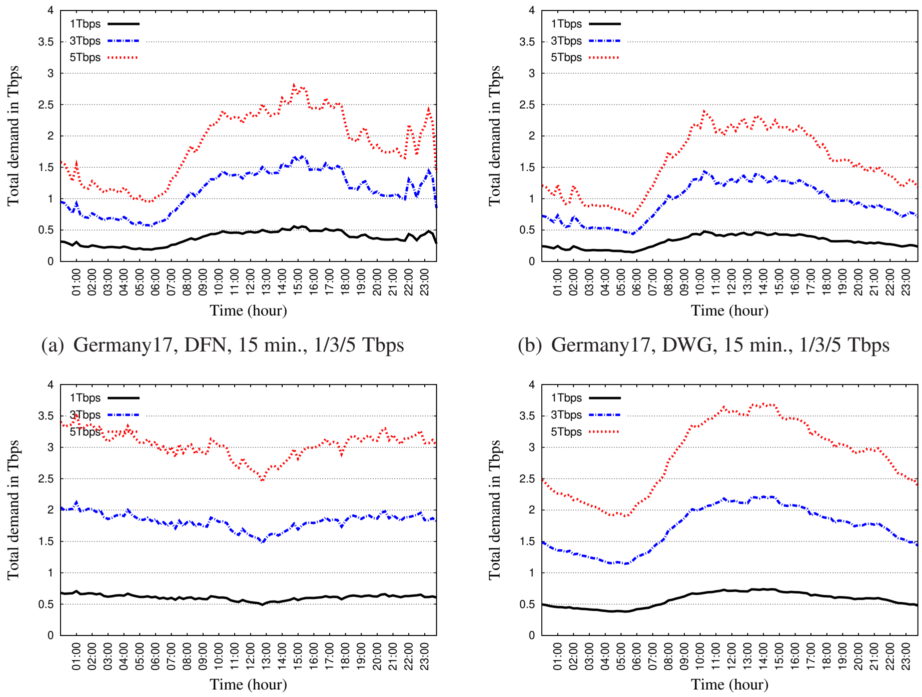

(a) Germany17, DFN, 15 min., 1/3/5 Tbps

(b) Germany17, DWG, 15 min., 1/3/5 Tbps

(c) Abilene, 15 min., 1/3/5 Tbps

(d) Géant, 15 min., 1/3/5 Tbps

Figure 4: Total demand over a day in Tbps for the scenarios with 1/3/5 Tbps maximum total demand. The traffic shows different levels of dynamics. The Germany17 (a), (b) and Géant demand (d) shows a typical day/night pattern with more short-range dynamics for Germany17. The Abilene demand (c) is quite flat over the day.

Observatory [45] and is available under [48]. The original data of 5-minute time granularity has been aggregated to similar time intervals as for the DFN measurements (the whole year could not be covered due to data unavailability - we considered time intervals of 2 weeks between 2004-05-01 and 2004-07-24 instead). The traffic with 15-minute intervals over 2005-08-05 is shown in Fig. 4 (c).

Another set of traffic matrices is available for the Géant network [41] It was collected using Netflow statistics and BGP Routing Information Base (RIB). We would like to acknowledge Steve Uhlig for granting us access to nonanonymized data. We keep the dates of the measurements and absolute demand values anonymized for the Géant traffic matrices. Please note that the total demand values 1 Tbps, 3 Tbps and 5 Tbps are scaled values. We do not report the scaling factors in our work. We considered traffic matrices with the original time granularity of 15 minutes over a day (Fig. 4 (d)), and time intervals of a day over a month.

The traffic values for all networks have been converted to Mbps and scaled. To get undirected demands between nodes i and j we considered the maximum of the two corresponding directed demands.

As shown in Fig. 3 (a), the DFN matrices have a centralized structure with a large demand emanating from Frankfurt, which is a large entry point for cross-atlantic traffic, see [8]. 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 Germany17 with demand matrices generated using the Dwivedi-Wagner (DWG) model [13] based on population statistics. The resulting demands are much less centralized. Fig. 3 highlights the differences between the DFN and DWG demands; the area of each node is proportional to its emanating demand. This model distinguishes between data, voice, and video traffic and computes demand values between two cities according to their distance and their number of inhabitants, employees, or households depending on the traffic class. 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

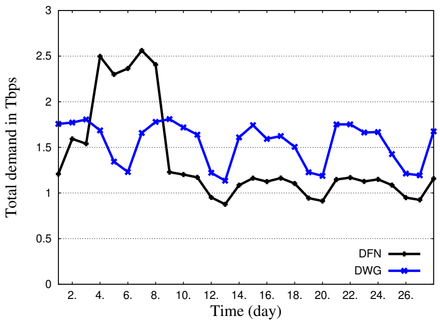

(a) Germany17, 1 day, 3 Tbps

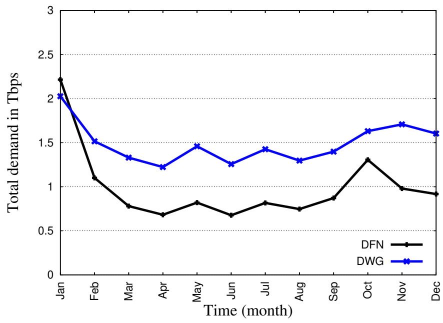

(b) Germany17, 1 month, 3 Tbps

Figure 5: Total demand over time in Tbps for the scenario with 3 Tbps maximum total demand, every day of February 2005, and every month of 2004. The DWG matrices show the expected behavior over time with peaks during the week, low traffic on weekends (a), and a slightly rising demand over the year (b). In contrast, the DFN measurements exhibit peaks caused by single demands.

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 $[0, 1]$ normalizing the maximum DFN demand for every (i, j). It hence rules out the domination effects caused by single demands in the measurements. Fig. 5 illustrates this effect for the daily total demand values over a month (a), and monthly total demand values over a year (b). The topologies and traffic data are summarized in Table 1.

6 Results

This section describes our computational results on the data presented in the previous section. We first explain the essence of our results using the Germany17 network and the 96 DFN traffic matrices given for every 15 minutes of 24 hours, and discuss these results in detail. Afterwards, we illustrate that most of these results are invariant against changes of

- the ratio between demand values and the capacity granularity (demands scaled to 1, 3 and 5 Tbps maximum total demand),
- the time scale of the demands (every 15 minutes over one day, every day over a month, and every month over a year),
- the network (Germany17, Abilene, and Géant), and

Table 1: Network topologies and traffic

Network	Nodes	Links	Time granularity	Reference
Abilene	12	15	every 15 min. of 2005-08-05, every day of 2004-05, every two weeks between 2004-5-01 and 2004-07-24	[48]
Géant	22	36	every 15 min. over a day, every day over a month	[41]
Germany17	17	26	every 15 min. of 2005-02-15, every day of 2005-02, every month of 2004 using DFN-measurements and Dwivedi-Wagner model	[3]

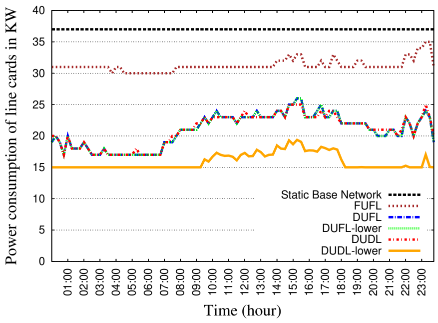

(a) 1 Tbps

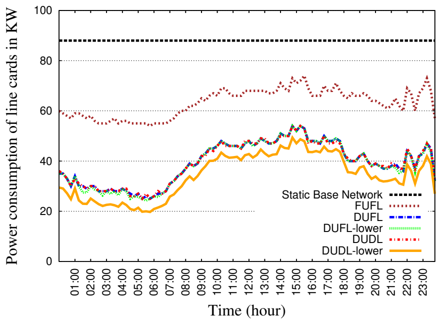

(b) 3 Tbps

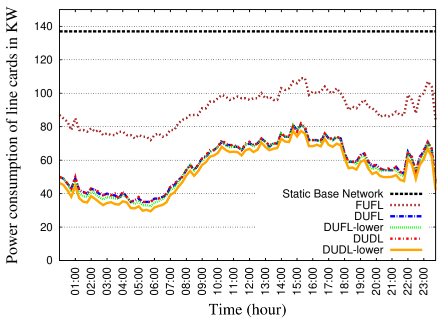

(c) 5 Tbps

Figure 6: The figures show the power consumption in kilowatt with the three strategies on the Germany17 network, DFN traffic, 1/3/5 Tbps, every 15 minutes on February 15, 2005. The difference between FUFL and DUFL is much larger than the additional benefit of DUDL. Note that the y-axes are not identical.

- the structure and spatial distribution of the demand matrix (DFN measurements or Dwiwedi-Wagner model using population statistics).

while we report on some influence of the regional distribution of traffic (measurements vs. population statistics) on the success of FUFL. All occurring MILPs have been solved using CPLEX 12.1 [18] as a black-box solver with a time-limit of one hour on a 64-bit Intel 3.00 GHz CPU with 8 GB main memory.

6.1 Results for the Germany17 network with DFN matrices over one day.

We could compute an energy-minimal solution within seconds or minutes for most of the DUFL scenarios on the Germany17 network with the 15-minute DFN matrices for a day. 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), 6%–25% (3 Tbps), and 3%–15% (5 Tbps). A higher relative optimality gap for DUDL can be observed for the 1 Tbps maximum total demand than for 5 Tbps. This is the usual behavior when changing the ratio between the demand values and the capacity granularities. All comparisons of the three strategies are 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 there no dual bounds for FUFL since no optimization is performed with this approach. For almost all DUFL runs dual bounds and solution values are identical which refers to energy optimal solutions. Whenever reporting on power values we consider the total power consumptions of all active line-cards assuming a value of 500 W for every single line card as explained in Section 5.

Figures Fig. 6(a)–Fig. 6(c) illustrate the optimal power consumptions obtained with FUFL as well as the power consumptions and dual bounds for DUFL and DUDL for each of the three demand scalings.

All the proposed approaches make use of the dynamics of traffic and follow the total demand curve (compare with Fig. 4(a)). The energy savings achieved with DUFL and DUDL are nearly identical and much larger than with FUFL. In other words, additional possibility of 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. The constant periods of power consumption in Fig. 6(a) are due to the minimal number of line cards that are needed to maintain IP connectivity.

More precisely, in the 1 Tbps scenario, the line cards of the network consume 0.89/0.75/0.50/0.38 MWh over the day for Base/FUFL/DUFL/DUDL, respectively. The corresponding values for 3 Tbps and 5 Tbps are 2.11/1.52/0.94/0.82 MWh and 3.29/2.15/1.36/1.28 MWh. Notice that these values correspond to the power consumption of a single day. Accumulating these values over a year (multiplying by 365) results in an energy consumption of 1201/785/496/467 MWh over the year in the 5 Tbps scenario, which is verified by our calculations using matrices for every month over a year.

Although we focus on the comparison of the FUFL, DUFL and DUDL approaches, it may be interesting from the practical perspective to evaluate the energy savings against the static base network. We observe that that already the easy-to-realize FUFL saves substantial energy compared to our base network. In the 3 Tbps scenario FUFL reduces the power of the active line cards in low-demand hours by up to 38% at 05:30 am (72% for DUFL and 77% for DUDL). Even in a high-load scenario the savings are significant (17%, 39%, and 44% for FUFL, DUFL and DUDL at 02:45 pm, respectively). One should not overestimate the total energy savings compared to the static base network, since it may be overprovisioned (peak demands typically do not occur simultaneously).

Considering the power consumption at 05:30 am and 02:45 pm for a maximum total demand of 3 Tbps, 25% of power for FUFL, 55% for DUFL, and 59% for DUDL can be saved in the early morning compared to the peak hour.

Fig. 7 shows the virtual topologies corresponding to the base network (a) and the considered approaches ((b) FUFL, (c) DUFL, (d) DUDL) for a low-demand hour in the 5 Tbps scenario (at 05:30 am, see Fig. 6). Although FUFL allows to reduce the link capacity (number of active line cards), the virtual topology remains nearly the same because existing virtual links have to be maintained even for small amount of traffic. In contrast, the virtual topology changes significantly with DUFL and DUDL. Moreover, lightpaths in DUFL and DUDL are highly utilized, as opposed to FUFL, see Fig. 8. Remember that DUFL may only use virtual links that exist in the base network in contrast to DUDL which may set up new lightpaths (e.g. Hannover–Norden). Nevertheless the virtual topologies of DUFL and DUDL are very close to each other with a similar number of active line cards.

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 (see Fig. 7(a)). Since DUFL may use any virtual link from the base network, the DUFL solution is (close to) optimal and 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. Both DUFL and DUDL will find a tree network. These trees might differ, but they use the same number of line cards. Hence again DUDL cannot benefit compared to DUFL. For the success of DUFL it is also crucial that we allow flow splitting in the virtual domain, which lets DUFL fill up the established lightpaths to a high extent. This is illustrated in Fig. 8(a) and Fig. 8(b).

6.2 Varying input parameters.

To make sure that the observed results do not depend on specific assumptions on the network or the input data, we have varied several parameters:

We evaluated the performance of FUFL, DUFL, and DUDL for all possible parameter combinations of the network, the time scale, and the demand scalings. Since the observed results were basically consistent over all these scenarios, we show an example diagram for each of the variations. For the purpose of this subsection, we consider the DFN 15min-measurements over a day scaled to 3 Tbps on the Germany17

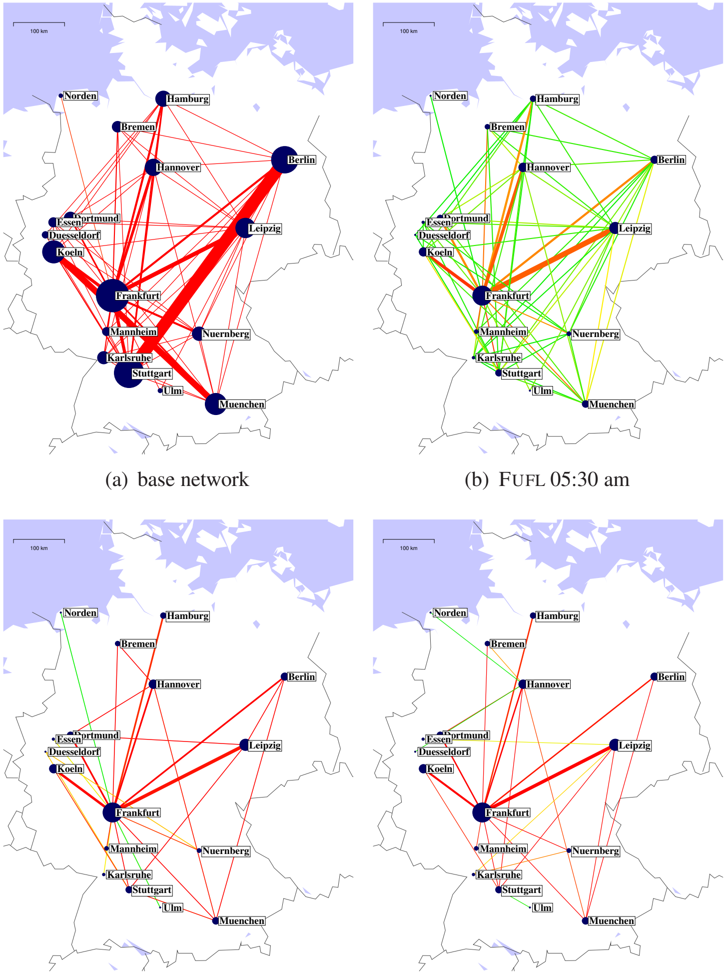

(a) base network (b) FUFL 05:30 am

(c) DUFL 05:30 am (d) DUDL 05:30 am

Figure 7: Virtual topologies of the base network and of the computed solutions with FUFL, DUFL, and DUDL at 05:30 am for the Germany17 network with DFN measurements, maximum total demand 5 Tbps. The color of a link corresponds to the load (high utilization (red) to low utilization (green)). The width of a link refers to its capacity. The size of a node represents its emanating demand.

network as the reference scenario, which is illustrated in Fig. 6(b). In the previous section, we have already shown that our results are basically invariant against variations of the ratio between demand and capacity granularities, compare Fig. 6(a), Fig. 6(b) and Fig. 6(c).

Varying the time scale Secondly, we varied the time scale, which changed the structure of the demand variations over time. For the DFN measurements on the Germany17 network, we not only considered the 15-minute demand matrices over a day, but also aggregated measurements for every day over a month, and for every month over a year, as explained in detail in Section 5. Fig. 9 shows the power consumption over time with FUFL, DUFL, and DUDL for the latter two time scales on the Germany17 network with DFN demands. One can see that DUFL and DUDL are very close to each other and can save much more energy than FUFL also on these time scales.

Varying the network Thirdly, we varied the network topology by using the traffic measurements on the Géant and Abilene networks described in Section 5. Fig. 10 shows the power consumption curves over time for these two networks. Again, the difference between the three approaches is similar to the results

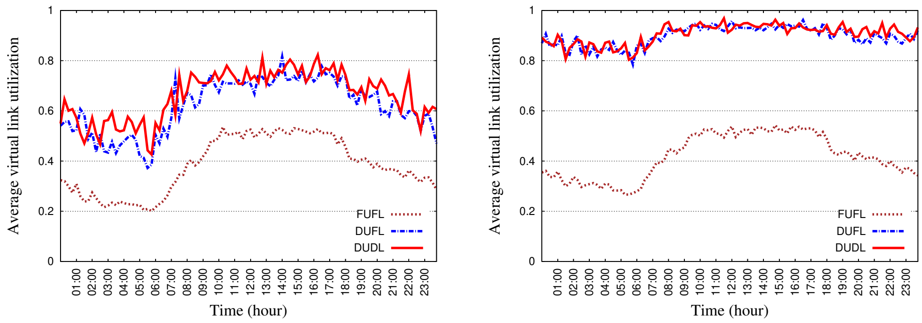

(a) 1 Tbps (b) 5 Tbps

Figure 8: Average virtual link utilization over all active virtual links, Germany17, DFN, every 15 minutes. DUDL and DUFL achieve high lightpath utilization in contrast to FUFL.

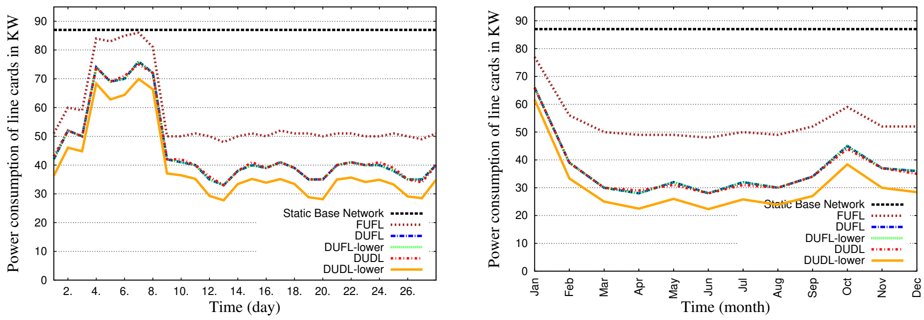

(a) Every day over a month (b) Every month over a year

Figure 9: The results are independent of the time scale. The figures show the power consumption over time with FUFL, DUFL, and DUDL on the Germany17 network, DFN traffic, 3 Tbps, for a month and a year. The meaning of the curves is the same as in the 15-minute figure Fig. 6(b).

observed on the other networks. Notice that on the Abilene network, all obtained network topologies are optimal.

Varying the demand pattern Eventually, we varied the structure of the demands by using demands generated with the Dwiwedi-Wagner (DWG) model instead of the DFN measurements. Fig. 11 shows the power consumption over time with the DWG demands on the Germany17 network. It turns out that the essential result does not change: The benefit of DUFL compared to FUFL is very large, and DUFL and DUDL differ only marginally. The plotted dual bounds show that the network topologies obtained in each time slot are close to be optimal. In other words, we proved that this effect is due to structural differences of the considered rerouting concepts and not the result of a heuristic solution procedure.

On the other hand, it can be seen that the energy reduction compared to the base network using the FUFL approach is marginal, which is in contrast to the results for the DFN-demands (compare Fig. 6(b) with Fig. 11). Recall that the DFN-traffic is very centralized with a large concentration of demand in Frankfurt, while the DWG-demands are more evenly distributed (compare Fig. 3(a) with Fig. 3(b)). This has mainly two consequences. First, the line-cards are more evenly distributed among the nodes in the 3 Tbps DWG solution with a minimum of 2 line cards (Norden) and a maximum number of 15 line cards (Frankfurt, Leipzig, and Muenchen) compared to a minimum and maximum of 1 (Norden and Ulm) and 32 (Frankfurt) in the 3 Tbps DFN base network solution. And secondly, the number of used physical paths is larger in the DWG solution compared to the DFN solution with fewer parallel lightpaths. In the DWG base network 87 physical paths are in use with a maximum number of 2 lightpaths using the same physical path. The DFN solution uses a total of 50 paths with a maximum of 14 channels on the same path. Note that total number of channels is almost the same in both scenarios (88 for DFN and 89 for DWG) because the sum of all demands is identical (3 Tbps). Since FUFL, as explained in Section 4, may switch off line cards only in the presence of parallel channels, energy can only slightly be reduced with DWG demands. Here the spacial distribution of the traffic influences the impact of FUFL approach on energy saving. In contrast, we do not observe a difference for DUFL and DUDL.

7 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. Our extensive computational study strongly suggests that these results are independent of the ratio between the demand and capacity granularities, the demand structure, the time scale, and the network topology. Already a simple monitoring of the traffic and downscaling the line card use accordingly (FUFL), may bring substantial savings in energy cost. We observed that the impact of this latter simple strategy depends on the regional distribution of the demand. With the same total traffic and the same dynamics over time, less energy can be saved with FUFL having (non-realistic) evenly distributed demands.

Traffic engineering and reconfiguring the IP routing is nowadays part of the daily business of network operators. Our work indicates that energy aspects should be considered in the operational phase helping to save OPEX. It should also motivate equipment vendors to provide network elements with convenient and fast functionality to be switched on and off.

We focused on a very small subset of energy-consuming network equipment in telecommunication networks. We considered only backbone topologies neglecting metro and access networks as well as equipment of end-users. We also restricted our attention to the equipment for network links. Assuming a power consumption of 3 kW for a single shelf without cards, the considered line cards are responsible for only roughly 25% (5 Tbps scenario) of the energy consumed by IP routers. However, we showed that the energy saved only by using flexible line card equipment and only in the backbone already amounts to almost 800 MWh over a year.

Finally we remark that we see some potential also in reducing the energy consumption of node equipment. Router chassis of smaller capacities may be used if the dynamics of IP traffic allow to switch off many active line cards at a certain node. Obviously, ISPs cannot dynamically exchange the IP router equip-

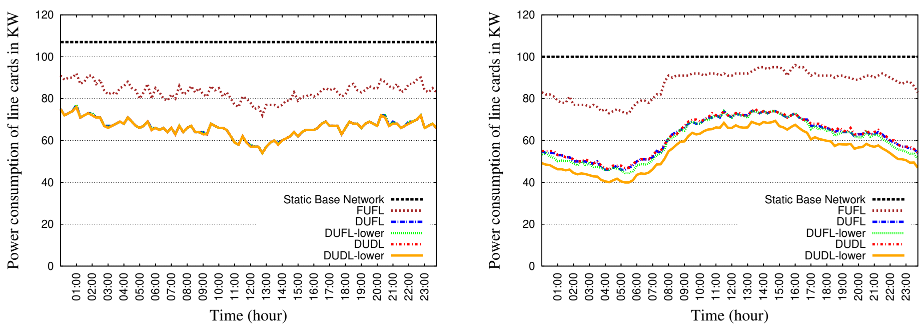

(a) Abilene network (b) Géant network

Figure 10: The results are independent of the network topology. The figures show the power consumption over time with the three approaches on the Géant and Abilene network with measured traffic demands, scaled to 3 Tbps maximum total demand, every 15 minutes of a day.

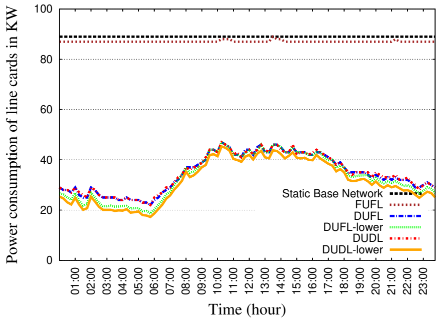

Figure 11: Power consumption over time with FUFL, DUFL, and DUDL on the Germany17 network, DWG demands, 3 Tbps, 15 minutes over one day. The results are basically independent of the demand structure. But FUFL fails to reduce the number of active line cards.

ment they installed. But in the context of multi-chassis configurations it might be possible to switch off at least single elements of IP routers (single shelves, slot cards, etc.).

References

[1] Cisco CRS-1 Production Brochure. `http://www.cisco.com/en/US/prod/collateral/routers/ps5763/prod_brochure0900aecd800f8118.pdf`, October 2008.

[2] Deutsche Forschungsnetz Verein. `http://www.dfn.de`, 2010.

[3] NOBEL – Next generation Optical networks for Broadband European Leadership. `http://www.ist-nobel.org`, 2004–2007.

[4] A. Altin, H. Yaman, and M. C. Pinar. The Robust Network Loading Problem under Polyhedral Demand Uncertainty: Formulation, Polyhedral Analysis, and Computations. *INFORMS Journal of Computing*, Articles in Advance, 2010.

[5] B. G. Bathula, M. Alresheedi, and J. M. H Elmirghani. Energy efficient architectures for optical networks. In *Proc. of London Communication Symposium, London, UK*, 2009.

[6] P. Belotti, A. Capone, G. Carello, and F. Malucelli. Multi-layer mpls network design: The impact of statistical multiplexing. *Computer Networks*, 52(6):1291–1307, 2008. ISSN 1389-1286.

[7] D. Bienstock and O. Günlük. Capacitated network design – polyhedral structure and computation. *Informs Journal on Computing*, 8:243–259, 1996.

[8] A. Bley, U. Menne, R. Klähne, C. Raack, and R. Wessäly. Multi-layer network design – a model-based optimization approach. In *Proceedings of the PGTS 2008, Berlin, Germany*, pages 107–116. Polish-German Teletraffic Symposium, 2008.

[9] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright. Power awareness in network design and routing. In *Proceedings of the IEEE INFOCOM 2008, Phoenix, USA*, pages 1130–1138, April 2008.

[10] L. Chiaraviglio, M. Mellia, and F. Neri. Energy-aware backbone networks: a case study. In *Workshop on Green Communications 2009, Dresden, Germany*, June 2009.

[11] L. Chiaraviglio, M. Mellia, and F. Neri. Reducing power consumption in backbone networks. In *Proceedings of the IEEE ICC 2009, Dresden, Germany*, June 2009.

[12] G. Dahl, A. Martin, and M. Stoer. Routing through virtual paths in layered telecommunication networks. *Operations Research*, 47(5):693–702, 1999.

[13] A. Dwivedi and R. E. Wagner. Traffic model for USA long distance optimal network. In *Proceedings of the Optical Fiber Communication Conference*, pages 156–158, March 2000.

[14] B. Fortz and M. Poss. An improved benders decomposition applied to a multi-layer network design problem. *Operations Research Letters*, 37:359–364, 2009. doi: 10.1016/j.orl.2009.05.007.

[15] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, San Francisco, New York, 1979.

[16] M. Gupta and S. Singh. Greening of the internet. In *Proceedings of ACM SIGCOMM'03, Karlsruhe, Germany*, August 2003.

[17] 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*, 7(9):814–833, 2008.

[18] IBM – ILOG. CPLEX. `http://www.ilog.com/products/cplex/`, 2009.

[19] F. Idzikowski. Power consumption of network elements in IP over WDM networks. Technical Report TKN-09-006, Technical University of Berlin, Telecommunication Networks Group, July 2009.

[20] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz. Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios. In *Proceedings of the 14th Conference on Optical Networks Design and Modeling (ONDM 2010), Kyoto, Japan*, February 2010.

[21] A. M. C. A. Koster and A. Zymolka. Minimum converter wavelength assignment in all-optical networks. In *Proceedings of the 8th Conference on Optical Network Design & Modelling (ONDM)*, pages 517–535, 2004.

[22] A. M. C. A. Koster, S. Orlowski, C. Raack, G. Baier, and T. Engel. *Single-layer Cuts for Multi-layer Network Design Problems*, volume 44, chapter 1, pages 1–23. Springer Verlag, February 2008. Selected proceedings of the 9th INFORMS Telecommunications Conference.

[23] A. M. C. A. Koster, S. Orlowski, C. Raack, G. Baier, T. Engel, and P. Belotti. *Branch-and-cut techniques for solving realistic two-layer network design problems*, chapter 3, pages 95–118. Springer Berlin Heidelberg, December 2009.

[24] E. Kubilinskas. *Design of Multi-layer Telecommunication Networks*. PhD thesis, Lund University, 2008.

[25] X. Liu. The role of stochastic programming in communication network design. *Comput. Oper. Res.*, 32(9):2329–2349, 2005. ISSN 0305-0548.

[26] T. L. Magnanti, P. Mirchandani, and R. Vachani. Modelling and solving the two-facility capacitated network loading problem. *Operations Research*, 43:142–157, 1995.

[27] S. Mattia. The robust network loading problem with dynamic routing. Technical Report 3, Università di Roma la Sapienza, 2010.

[28] S. Orlowski. *Optimal Design of Survivable Multi-layer Telecommunication Networks*. PhD thesis, Technische Universität Berlin, May 2009.

[29] F. Ortega and L. A. Wolsey. A branch-and-cut algorithm for the single-commodity, uncapacitated, fixed-charge network flow problem. *Networks*, 41:143–158, 2003.

[30] E. Palkopoulou, D. A. Schupke, and T. Bauschert. Energy efficiency and CAPEX minimization for backbone network planning: Is there a tradeoff? In *Proc. of the ANTS, New Delhi, India*, 2009.

[31] F. Pascali, M.G. Scutellà, and à. Chance constrained network design. In *Proceedings of the Fourth International Network Optimization Conference (INOC 2009), Spa, Belgium*, 2009.

[32] M. Pióro and D. Medhi. *Routing, Flow, and Capacity Design in Communication and Computer Networks*. Morgan Kaufmann Publishers, 2004.

[33] B. Puype, D. Colle, M. Pickavet, and P. Demeester. Energy efficient multilayer traffic engineering. In *Proceedings of the ECOC 2009, Vienna, Austria*, September 2009.

[34] B. Puype, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester. Power reduction techniques in multilayer traffic engineering. In *ICTON Conference*, June 2009.

[35] S. Raghavan and D. Stanojević. WDM optical design using branch-and-price. Working paper, Robert H. Smith School of Business, University of Maryland, April 2007.

[36] J. C. C. Restrepo, C. G. Gruber, and C. M. Machuca. Energy profile aware routing. In *Workshop on Green Communications, Dresden, Germany*, 2009.

[37] G. Shen and R. S. Tucker. Energy-minimized design for IP over WDM networks. *Journal of Optical Communications and Networking*, 1(1):176–186, 2009.

[38] W. Shen, Y. Tsukishima K. Yamada, and M. Jinno. Power-efficient multi-layer traffic networking: Design and evaluation. In *Proc. of the ONDM, Kyoto, Japan*, 2010.

[39] A. Silvestri, A. Valenti, S. Pompei, F. Matera, and A. Cianfrani. Wavelength path optimization in optical transport networks for energy saving. In *Proc. of the ICTON, Azores, Portugal*, 2009.

[40] S. E. Terblanche, R. Wessäly, and J. M. Hattingh. Solution strategies for the multi-hour network design problem. In *Proceedings of the Third International Network Optimization Conference (INOC 2007), Spa, Belgium*, April 2007.

[41] S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon. Providing public intradomain traffic matrices to the research community. *ACM SIGCOMM Computer Communication Review*, 36(1):83–86, 2006.

[42] United States Department of Energy. International energy outlook 2008, report doe/eia-0484(2008). `http://www.eia.doe.gov/oiaf/ieo/pdf/0484(2008).pdf`, September 2008.

[43] N. Vasić and D. Kostić. Energy-aware traffic engineering. Technical Report NSL-REPORT-2008-004, School of Computer and Communication Sciences, EPFL, Switzerland, 2008.

[44] N. Vasić and D. Kostić. Energy-aware traffic engineering. In *Proc. of 1st International Conference on Energy-Efficient Computing and Networking, Passau, Germany*, 2010.

[45] H. Wang, H. Xie, L. Qiu, Y. R. Yang, Y. Zhang, and A. Greenberg. Cope: Traffic engineering in dynamic networks. *ACM SIGCOMM Computer Communication Review*, 36(4):99–110, 2006.

[46] Y. Wu, L. Chiaraviglio, M. Mellia, and F. Neri. Power-aware routing and wavelength assignment in optical networks. In *Proceedings of the ECOC 2009, Vienna, Austria*, September 2009.

[47] E. Yetginer and G. N. Rouskas. Power efficient traffic grooming in optical wdm networks. In *Proc. of the GLOBECOM, Honolulu, USA*, 2009.

[48] Y. Zhang. 6 months of abilene traffic matrices. `http://www.cs.utexas.edu/~yzhang/research/AbileneTM/`, August 2009.

[49] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg. Fast accurate computation of large-scale ip traffic matrices from link loads. *ACM SIGMETRICS Performance Evaluation Review*, 31(1):206–217, 2003.