Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ANDREAS BETKER[1], INKEN GAMRATH, DIRK KOSIANKOWSKI[1],
CHRISTOPH LANGE[1], HEIKO LEHMANN[1], FRANK PFEUFFER,
FELIX SIMON, AXEL WERNER

Comprehensive Topology and Traffic Model of a Nationwide Telecommunication Network

[1]Deutsche Telekom AG, Telekom Innovation Laboratories, Winterfeldtstr. 21, 10781 Berlin, Germany

Supported by the German Federal Ministry for Economic Affairs and Energy (BMWi) within the IT2Green programme (www.it2green.de) and the German Federal Ministry of Education and Research (BMBF) within the EUREKA project SASER.

ZIB Report 14-36 (September 2014)

Herausgegeben vom
Konrad-Zuse-Zentrum für Informationstechnik Berlin
Takustraße 7
D-14195 Berlin-Dahlem

Telefon: 030-84185-0
Telefax: 030-84185-125

e-mail: bibliothek@zib.de
URL: http://www.zib.de

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

Comprehensive Topology and Traffic Model of a Nationwide Telecommunication Network

Andreas Betker[†], Inken Gamrath[‡], Dirk Kosiankowski[†], Christoph Lange[†], Heiko Lehmann[†], Frank Pfeuffer[‡], Felix Simon[‡], Axel Werner[‡]

24 Sep 2014

Abstract

As a basis for meaningful simulation and optimization efforts with regard to traffic engineering or energy consumption in telecommunication networks, suitable models are indispensable. This concerns not only realistic network topologies, but also models for the geographical distribution and the temporal dynamics of traffic, as well as the assumptions on network components and technology. This paper derives such a model from the practice of a large national carrier.

Applying the network and traffic model, we demonstrate its use by presenting various optimization cases related to energy-efficient telecommunication. Here, we focus on load-adaptivity by employing sleep modes to the network hardware, where several constraints on the reconfigurability of the network over time are considered.

1 Introduction

With the evolution from a simple voice transmitting device to a carrier of a variety of data transmission services with ever increasing traffic demands, telecommunication networks have become subject to incessant reconstruction to keep up with the growing requirements. Since the rising acquisition, installation and operation cost of network equipment has become one of the major concerns to the networks' operators, researchers study how to design and operate networks in a cost-efficient fashion; they rely on the availability of up-to-date network infrastructure, architecture and traffic data to base their studies on. To mitigate some shortcomings of existing data sets, we present a reference layer-2/layer-3 over WDM network mimicking the features of the telecommunication network of a nationwide network operator. Additionally we describe a traffic model comprising typical traffic variations over a day for the reference network. Based on this data we show a number of exemplary computations concerning power- and energy-optimal network configurations.

There are several telecommunication network topology and traffic models published in the literature, see, for instance, [11, 5, 6, 23, 18]. In the traffic model of Dwivedi and Wagner [11], population density statistics and geographical distances are used to estimate the average traffic demands between major cities. It distinguishes between traffic from three application areas, each

[†]A. Betker, D. Kosiankowski, C. Lange, H. Lehmann, Deutsche Telekom AG, Telekom Innovation Laboratories, Winterfeldtstr. 21, 10781 Berlin, Germany, {andreas.betker, dirk.kosiankowski, christoph.lange, h-lehmann}@telekom.de

[‡]I. Gamrath, F. Pfeuffer, F. Simon, A. Werner, Zuse Institute Berlin (ZIB), Takustr. 7, 14195 Berlin, Germany, {inken.gamrath, pfeuffer, simon, werner}@zib.de

exhibiting its own characteristic growth rate per year and dependence on population composition, consumer habits and distance. The model is applied in [11] to a 46-node network (USA), and in [5] and [6] to a 17-node (Germany), a 28-node (Europe) and a 14-node (USA) network. None of the resulting data sets, however, feature dynamic traffic variations over shorter time spans like, for instance, a day.

In [23] temporally resolved traffic measurements are published for a 23-node research network (Europe), in [18] measurements for a 50-node research network (Germany) are used, and further network topologies with dynamic traffic data are available in the SNDlib [19]. These measurement-based data sets have the drawback that they do not distinguish between traffic from different application types. Also, their origin usually lies in public research networks, which may exhibit other traffic characteristics than commercial networks.

The cited network models lack a number of further features. They represent network structures that are mostly outdated by now: Operators currently consolidate their networks towards using fewer but larger core network sites and ports supporting higher bit rates, to increase cost-efficiency. Likewise, all the mentioned traffic demand sets are based on forecasts or measurements from 2005 or before. Thus, the traffic volumes are too low with respect to current observations: global IP traffic has increased eightfold from 2007 to 2011 alone [10]. Furthermore, many models offer only static traffic data whereas for assessing load-adaptive networking concepts dynamic traffic matrices are needed that model the fluctuating traffic over time. Finally, none of the above models contains information on the topology or the traffic of the aggregation sections, but rather concentrate on the core section of the network.

In this article, we present a reference model that addresses these missing issues. The network and traffic model is described in Section 2 and features a topology designed from the contemporary perspective of a commercial network operator. Apart from a core network, it includes a description of the aggregation network topology and dynamic traffic matrices for the core and aggregation parts. The latter are based on traffic forecasts for 2012 as well as measurements of a network operator and distinguish between four types of traffic sources. In Section 3 we present some exemplary applications of the network and traffic model. The computations concentrate on energy efficiency in load-adaptive networks, which is an important recent trend in network design literature, cf. [16], but obviously, the described model is also valuable for all kinds of further applications concerning dynamic networks. Section 4 summarizes the main contributions of the paper.

The detailed network topology together with the traffic data can be requested from the authors (email: `andreas.betker@telekom.de`).

2 Telecommunication Network Model

In this section we introduce the network model, both in terms of network topology as well as dynamic traffic characteristics.

2.1 General

The hierarchical structure of fixed-line communication networks typically consists of access, aggregation and core sections (see Figure 1).

Starting with the network end points, access networks gather the data from the end users. They are designed as star-/tree-like structures without protection. The traffic is then fed into the aggregation networks, which exhibit a star-/tree-like or weakly meshed structure with redundant connections. The aggregation network nodes are connected to the core network, which is typically highly or fully meshed. The aggregation part of the network consist of several independent

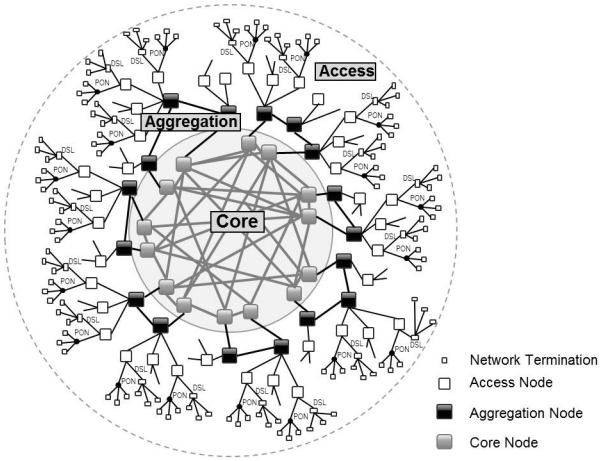

Figure 1: Telecommunication network topologies: Star-/tree-like and meshed network parts.

sections, each covering a certain confined region with its infrastructure, whereas the core network interconnects the regional aggregation sections using a nationwide infrastructure. More detailed information on the network segments and the current structure of telecommunication networks can be found, for example, in [17].

Redundant paths between nodes in partly or fully meshed structures are a prerequisite for flexibly routing and rerouting traffic. For this reason the subsequently presented model embraces only the aggregation and the core part of the telecommunication network, with an optical transport network underneath; the access networks are excluded from the considerations throughout this article since no alternative paths are available in these tree-shaped networks and only per-link adaptation techniques for the bit rate are applicable.

2.2 Network Topology Model

In the following, we provide a brief overview of the reference network model. The basis of the network presented here is hypothetical, but closely modelled on a realistic nationwide German optical communication network, comprising the aggregation and the core sections as well as the supporting optical transport network.

The network consists of 918 nodes and is divided into the aggregation part with 900 nodes and the core part with 18 nodes. The 18 core network nodes are located in 9 core network cities; there are always two nodes per core network city in different locations for resilience purposes.

Every city with an Internet exchange point (IXP) is represented by a core network city in the network model. The IXPs are in Frankfurt, Hamburg, Berlin, Dusseldorf, Stuttgart, Munich and Nuremberg. The IXP capacity in every city splits equally among both core nodes of this city. Reflecting the high traffic volume generated and the regional relevance, in addition to the existing IXP cities also Hanover and Leipzig are modeled as core network cities.

The optical core architecture consists of 26 fiber links, 17 between core nodes of different core network cities and one between the two core nodes of each of the 9 core network cities. Figure 2 shows both the optical layer and the IP layer, the latter featuring a considerably higher link-density.

For the purposes of protection and resilience, the IP layer is divided into two parallel subnetworks. Both subnetworks are designed in such a way that each of them can handle the whole

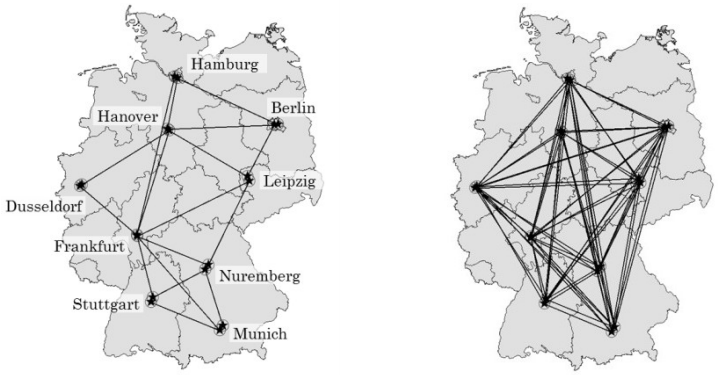

Figure 2: Core network: Optical layer (A) and IP layer (B).

traffic volume in case of disturbances or network failures: One node in every core city is allocated to the first subnetwork, the second node is part of the other subnetwork. Every subnetwork is fully meshed. Furthermore, in every core network city the node pair is connected. These design rules result in 81 links in the IP layer, 36 in every subnetwork and 9 internal core network cities' links.

All optical links are equipped with dense wavelength division multiplex (DWDM) systems. The core system's capacity is 8 Tbps (up to 80 wavelengths at 100 Gbps). Every core node consists of an IP router and an optical cross-connect (OXC) or a reconfigurable optical add-drop multiplexer (ROADM) on the optical layer. Each aggregation system's capacity is 400 Gbps (up to 40 wavelengths at 10 Gbps) and connects five aggregation nodes to both of the associated core nodes in a disjoint fashion. The add-drop capacity of a DWDM aggregation system is 800 Gbps (400 Gbps west and 400 Gbps east), so the average add-drop capacity of each of the five nodes within an aggregation system is 160 Gbps, 80 Gbps in every direction. Every aggregation node is equipped with a layer-2/3 switch which is linked to a (fixed) optical add-drop multiplexer (OADM). To amplify the optical signal on the fiber, every 80 km an optical line amplifier (OLA) is used. The link length between two connected nodes is assumed to be the linear distance multiplied by a factor of 1.2 , a value derived from existing networks. These systems are equipped with an adequate number of line cards and pluggable transceivers to meet the traffic requirements.

All optical network connections are 1+1 protected by a link-disjoint pair of paths, i. e., all data traffic is conducted independently over two disjoint link sequences and therefore is protected against optical link failures in the entire network.

The number of aggregation nodes corresponds to the number of Bitstream Access Points of Deutsche Telekom [22]. The selection of the aggregation nodes is based on a weighting of population density and takes into account a uniform distribution of nodes for nationwide coverage.

Every aggregation node is connected to an associated core network city via two optical fiber links, running along disjoint paths. The optical network implementation is based on so-called open-rings: These are optically transparent domains with topological ring structures. However, the ring is not closed in one single core network node, but in the core network city's second node instead. All optical channels are optically terminated with transponders in two terminal optical multiplexers (OMUX), one on each core node. The connection of an open-ring to both nodes of

a core network city is sketched in Figure 3.

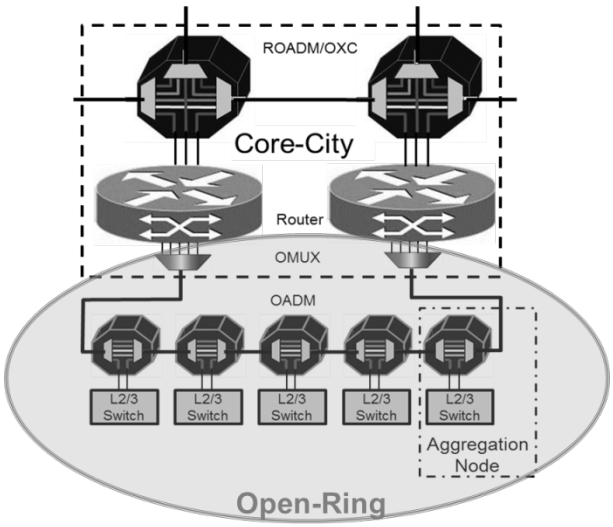

Figure 3: Connection of an open-ring aggregation structure to a core city.

Every open-ring structure consists of five aggregation nodes; in total, there are 180 of such structures, together comprising all 900 aggregation nodes. Figure 4 depicts the aggregation network with an enlarged part showing some aggregation structure details. The points marked with stars represent both core nodes in the core network city Berlin, the other points are aggregation nodes; the right part of the figure shows an example of an open-ring with five aggregation nodes in the region north east of Berlin.

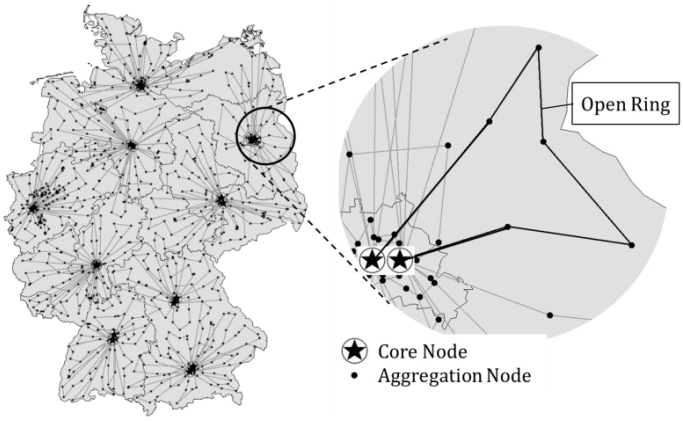

Figure 4: Aggregation network in Germany with enlarged part of the region north east of Berlin.

On a logical level every aggregation node is directly connected to both routers in the associated core city forming a star-like architecture. There is no direct logical connection among aggregation nodes themselves; the result is a hub and spoke network over an optical transparent open-ring.

The interconnection between aggregation and core network is made by using standardized interfaces. There is no transparent crossover between the domains for the following reasons:

- Operation, Administration and Maintenance (OAM) and performance monitoring on the domain border,
- possible use of different systems (of different vendors),
- grooming from lower to higher granular traffic between aggregation and core,
- breakup of wavelength blockades.

2.3 Network Traffic Model

The basis of the traffic model are forecasts of the Cisco Visual Networking Index (VNI) [10].The VNI is an ongoing effort to forecast and analyze the growth and use of IP networks worldwide. The VNI estimates different types of traffic in different countries in terms of average and busy hour traffic over a five-year period. For the purpose of the traffic modelling presented in this article, these forecasts were completed and refined by using internal forecasts on the basis of regional measurements of a big telecommunication provider. The resulting forecast for the total traffic in 2012 showed a peak data flow of approximately 10 Tbps in Germany (Figure 5).

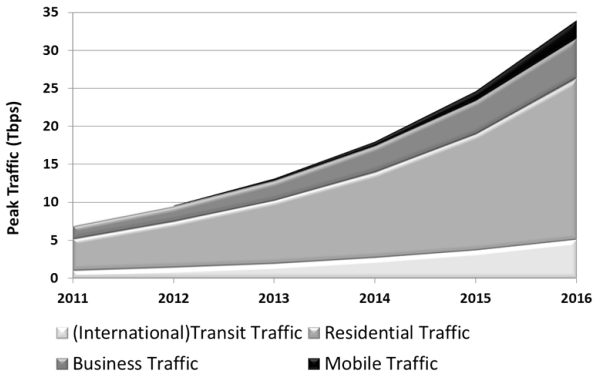

Figure 5: Peak traffic forecast for Germany.

Structurally, this traffic volume is generated from:

- residential traffic originating from approximately 36 million private households' telecommunication connections [9],
- business traffic from over 25 million business computers [9] in over 2 million companies [21] (approximately 20 % of the total traffic),
- mobile radio traffic from 62 million mobile users of radio networks [9] (approximately 1 % of the total traffic, but growing over-proportionally),
- core transit traffic between public Internet exchange points (IXP). This part of the total core traffic has been estimated in different non-public studies to be in a range of 10 – 20 %. The calculations in this article assume a portion of 15 %.

For the model, the following traffic assumptions are made:

- Every household with a telecommunication connection is a residential traffic source.
- Every office computer with internet connection is a business traffic source.
- Every mobile user is a mobile radio traffic source.
- Every user in its group utilizes the same bandwidth.
- Core transit traffic is part of the transit traffic which is neither originating nor terminating in Germany.

Based on this categorization, Figure 6 shows the predicted average bandwidth per traffic source during the overall traffic busy hour for residential, business and mobile subscribers as a function of time. Note that the overall busy hour is not necessarily at the same time as the

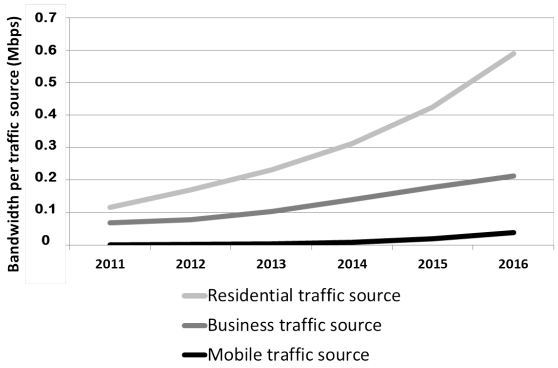

Figure 6: Busy hour bandwidth forecast per traffic source.

busy hour of the respective traffic type. Similarly the peak traffic times for the different types of traffic do not coincide in general, as shown below.

Every aggregation node is assigned to a particular region. The 900 aggregation regions cover the total area of Germany. All users inside this area are connected via the access network to dedicated aggregation nodes. So the generated traffic of every aggregation node depends on the number of users within its region and their individual traffic volume. Every aggregation node is associated with a core network node pair in a core network city; the working path leads to one core node, the protection path to the other core node located in the same city.

Figure 7 shows the regional traffic volume per area on an average working day, normalized to the peak traffic. The darker an area is shaded, the more traffic volume is generated there: It becomes obvious that in large cities and agglomeration areas major parts of the traffic volume are generated, whereas in rural regions the traffic demand is rather low.

Generally, there are two distribution modes of business, residential, and mobile traffic. The dominating traffic part is between user and IXP and the smaller part is the traffic between different individual users, so called peer-to-peer traffic. The traffic between users and IXP is divided into inner-aggregation traffic between the aggregation node and the corresponding core node and core traffic between this core node and the other core nodes. The distribution depends on the used IXP capacity. The peer-to-peer traffic is based on the regional user concentration. Traffic between nodes within an aggregation open-ring is always redirected over the core network

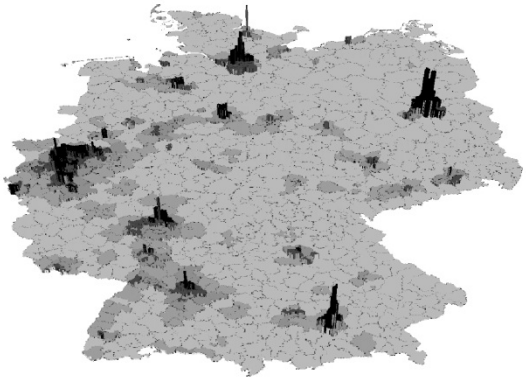

Figure 7: Normalized regional traffic volume per area.

node. Core transit traffic only occurs in the core network and its distribution also depends on the used IXP capacity.

This calculation results in 9 regional traffic matrices and one core traffic matrix. The regional matrices are the sums of residential, business, and mobile traffic between every aggregation node and its corresponding core network city. The core network traffic matrix contains the total traffic between all core network node pairs – consisting of residential, business, mobile, and international traffic. These matrices are the starting points for network optimization, see Section 3.

An important objective of this work is the modelling of temporally varying traffic flows, with respect to different traffic types. To this end, the following was taken into account:

- defining a flow curve for every traffic type over time,
- generating a corresponding traffic matrix for each point in time.

Several internal measurements within an operator's network, plus other regional non-public measurements returned for an average working day and an average weekend day different residential and business traffic flows, whereas the average mobile traffic flow shows negligible differences between the days of the week. While the weekly residential traffic peak is on the weekend, the business traffic peak is obviously on a working day.

To determine the flow curve for core transit traffic, data has been evaluated from [3] and [12]. The average traffic flow curve is the result of an observation of many requests over a longer period of time. Once again, the difference between the weekdays at the observed IXPs was negligible, and so the same flow curve is the basis for the calculations presented in this article.

Figure 8 illustrates the flow of all traffic parts on an average working and weekend day. The shapes of the traffic curves as functions of time have been obtained by non-public measurements and then scaled to match the peak-hour traffic values obtained from the VNI-based modelling. The residential traffic is the largest share of the total traffic, and the main differences between working and weekend days are found in the business traffic. The mobile radio traffic part today is marginal, but for the future above-average growth rates have been forecasted for this traffic type.

The total daily flow curve is the result of the sum of all traffic modes, plus 20 % reserve. The latter value is derived from the buffer ensuring the availability of the bandwidth guaranteed

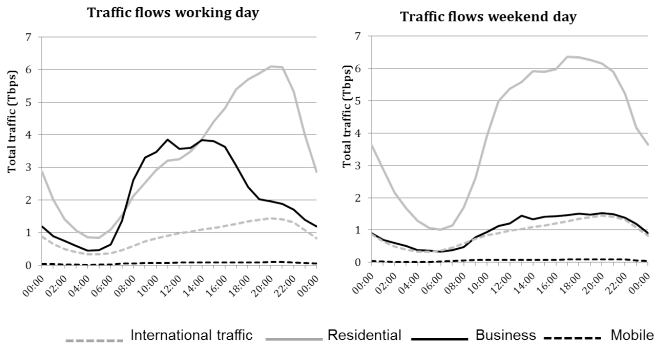

Figure 8: Traffic flows on different days of a week.

to business costumers and the general use of a maximum line utilization below 100 % (typically 80 %). Figure 9 illustrates the different daily traffic flows.

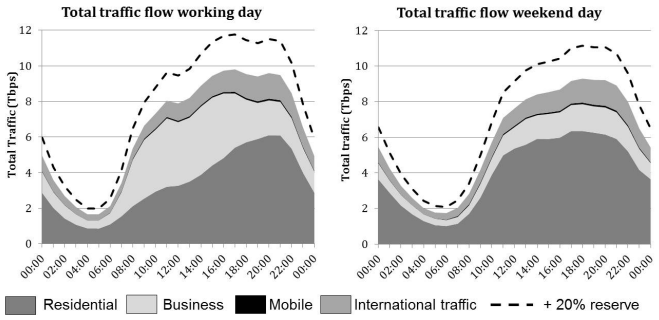

Figure 9: Exemplary traffic flow composition in a backbone network on a weekday and a weekend day.

The resulting curve is the basis for obtaining temporally varying traffic between every aggregation and core node. For each point in time, the residential, business and mobile traffic at a given aggregation node is obtained as the total respective traffic from the flow curve at the specified time, scaled with the relative number of users at the considered node. Traffic at a given core node is derived from the values at the corresponding aggregation nodes and the ratio between regional and long-distance traffic, plus the node's share of international traffic, depending on the dedicated traffic amount of the IXPs. These traffic values for each single node are used to create temporally varying traffic matrices for the core and aggregation network domains by distributing the node traffic among all pairs of nodes.

The traffic matrices can now be used for a variety of computations, such as estimating the power-saving potential due to load-adaptive network operation or deriving optimal switching levels to minimize network energy consumption. In such an application of the network and traffic model, further important side constraints may be taken into account to ensure the technical and operational realizability of the resulting traffic flows. The next section illustrates exemplary calculations and optimizations.

3 Application to Energy Efficiency Optimization

The network and traffic model described in the last section provides a sound basis for computations concerning dynamically changing networks. An important incentive for dynamic networks is to provide load-adaptivity and therefore energy efficiency during the operation of the network. This topic has been studied in a number of recent publications, see, for instance, [16, 20, 8, 17, 1]. In the following we use the data described in Section 2 in an exemplary way to assess the energy-saving potential in a nationwide network, taking a number of different requirements into account.

The model for the hardware architecture was taken from [15]. The assumed hardware is composed of an IP router and an optical cross-connect (OXC) on each node. The IP router consists of one or several line card shelves, which have to be interconnected by a fabric card shelf as soon as there is more than one line card shelf present at the node. Line card shelves are equipped with a suitable number of line cards, whose ports are connected to the optical cross-connect using transceivers generating long reach optical signals. The optical signals are then multiplexed by the optical cross-connect onto fibers and sent through the network. Data on specifications and power values for the assumed hardware was taken from [14] and [13]; we additionally assumed that all hardware that is not processing traffic can be powered off, and that hardware in a power-off state consumes no power.

We give a brief description of the mathematical model used for the computations. Let T be the set of time steps considered over the given time horizon; for our one-day, half-hourly discretized traffic model, we have $|T| = 48$. Energy-optimal routings in the backbone network for a fixed time $t \in T$ were computed using a two-layer, multicommodity flow integer program. The notation for parameters and variables is given in Table 1. We consider the commodities, i. e, pairs of nodes with positive traffic between these nodes, to be undirected and assume a given fixed ordering on the node set V to avoid complicated notation for pairs of nodes. The value d_k^t represents the demand for commodity $k = (a, b)$, i. e., the forecasted traffic between nodes a and b in time step t, and is obtained from the traffic data described in Section 2. The constants q_i^{A} count the line cards that are needed at node i to connect the aggregation network to the core and can be computed before solving the integer program according to the dimensioning of the aggregation network at time t (cf. below).

A solution to the integer program given below tells us for each time step t:

- the routing of traffic over IP links (f variables),
- the routing of light paths over optical connections (g variables),
- the number of power-consuming devices at core router sites (variables p, q, r, w),
- the number of channels established on IP links (y variables) as well as the number of transmitting fiber links (z variables).

$$\min \sum_{i \in V} (E_{\mathrm{P}}\, p_i + E_{\mathrm{L}}\, q_i + E_{\mathrm{S}}\, r_i + E_{\mathrm{F}}\, w_i + E_{\mathrm{O}}) + \sum_{e \in E} E_e\, z_e \tag{1}$$

$$\text{s.t.} \sum_{j:\{i,j\} \in L} (f_{i,j}^k - f_{j,i}^k) = \begin{cases} 1 & \text{if } k = \{i, b\},\ i < b, \\ -1 & \text{if } k = \{a, i\},\ a < i, \\ 0 & \text{otherwise} \end{cases} \qquad \forall i \in V,\ k \in K \tag{2}$$

$$\sum_{k\in K} d_k^t\,(f_{i,j}^k + f_{j,i}^k) \le C\,y_\ell \quad \forall \ell = \{i,j\} \in L \tag{3}$$

$$\sum_{v:\{u,v\}\in E} (g_{u,v}^\ell - g_{v,u}^\ell) = \begin{cases} 2\,y_\ell & \text{if } \ell = \{u,j\},\ u < j, \\ -2\,y_\ell & \text{if } \ell = \{i,u\},\ i < u, \\ 0 & \text{otherwise} \end{cases}$$
$$\forall u \in V,\ \ell \in L \tag{4}$$

$$g_{u,v}^\ell + g_{v,u}^\ell \le y_\ell \qquad \forall \{u,v\} \in E,\ \forall \ell \in L \tag{5}$$

$$\sum_{\ell\in L} (g_{u,v}^\ell + g_{v,u}^\ell) \le N\,z_e \qquad \forall e = \{u,v\} \in E \tag{6}$$

$$\sum_{j:\{i,j\}\in L} y_{\{i,j\}} = p_i \qquad \forall i \in V \tag{7}$$

$$p_i \le M_\mathrm{P}\,q_i \qquad \forall i \in V \tag{8}$$

$$q_i^\mathrm{A} + q_i \le M_\mathrm{L}\,r_i \qquad \forall i \in V \tag{9}$$

$$r_i \le (M_\mathrm{S} - {}^1\!/_{M_\mathrm{F}})\,w_i + 1 \qquad \forall i \in V \tag{10}$$

$$w_i \le M_\mathrm{F} \qquad \forall i \in V \tag{11}$$

$$f, g, y, z, p, q, r, w \ge 0, \text{ integer}, \tag{12}$$

Constraints (2) and (3) model the IP routing and make sure that enough light paths are available; constraints (4) and (5) model the routing of the light paths over optical links and ensure survivability of the optical routing in the case of an optical link failure, while constraints (6) provide for enough fiber capacity to carry all light paths. Constraints (7) to (11) take care that sufficient router capacity is available.

The objective (1) sums up the power consumption of the router hardware and OXCs at the nodes and OLAs on the fibers. For further details about the model and variants applicable to other conditions, we refer to [7].

For the optimization of the aggregation network no changes in the routing are possible. Hence, reductions in the power consumption are achieved solely by switching off parallel connections that are unused at low traffic times, similar to the FUFL scheme in [16]. This is due to the open-ring structure and the assumed hardware environment in the aggregation network; in principle these assumptions can be relaxed and the optimization model extended accordingly, which would then allow for further power consumption potential.

Table 1: Notation used for the integer program computing the backbone routing at time $t \in T$

Parameters	
V	set of core nodes
$E \subseteq \binom{V}{2}$	set of optical fiber links
$L \subseteq \binom{V}{2}$	set of IP links
$K \subseteq \binom{V}{2}$	set of point-to-point commodities with non-zero demand
$d_k^t > 0$	demand for commodity $k \in K$ at time t
C	bit rate capacity of a light path
N	maximum number of light paths per fiber
M_{P}	maximum number of ports per line card
M_{L}	maximum number of line cards per line card shelf
M_{S}	maximum number of line card shelves per fabric card shelf
M_{F}	maximum number of fabric card shelves per core IP router site
q_i^{A}	number of line cards for aggregation network connections at $i \in V$
E_{P}	power consumed by one port
E_{L}	power consumed by one line card (without ports)
E_{S}	power consumed by one line card shelf (without line cards)
E_{F}	power consumed by one fabric card shelf (without line card shelves)
E_{O}	overhead power consumption at a router site
E_e	power consumed along optical link $e \in E$ (due to OLAs)
Variables	
$f_{i,j}^k$	IP routing (flow w. r. t. commodity $k \in K$ over IP link $(i,j) \in L$)
$g_{u,v}^\ell$	optical routing (lightpaths provided for IP link $\ell \in L$ on optical link $(u,v) \in E$)
y_ℓ	number of necessary light paths on IP link $\ell \in L$
z_e	number of necessary fibers on optical link $e \in E$
p_i	number of activated ports at node $i \in V$
q_i	number of activated line cards at node $i \in V$
r_i	number of activated line card shelves at node $i \in V$
w_i	number of activated fabric card shelves at node $i \in V$

3.1 Energy-optimal network configurations

Optimal network configurations according to each traffic matrix yield lower bounds on the network power consumption in the respective time interval. Figure 10 shows this minimal power consumption over a working day and a weekend day for the core network. As a point of com-

Figure 10: Power consumption of load-independent configuration (dashed), optimal network configurations half-hourly (solid), and optimal network configurations with 6 reconfigurations per day (dash-dotted), each for the traffic (gray) over one working day (left) and one weekend day (right).

parison we also show a load-independent network configuration that is able to route all demands of working and weekend day and is also optimized with respect to power consumption. It was computed by using the integer program given above with Equation (3) being replaced by

$$\sum_{k \in K} d_k^t \, (f_{i,j}^k + f_{j,i}^k) \leq C \, y_\ell \quad \forall \ell = \{i,j\} \in L, \, t \in T \tag{13}$$

to ensure a static routing covering the demands d_k^t at *all* time points $t \in T$.

Load-independent network operation consumes about 977 kW, which amounts to 23.4 MWh per day. A load-adaptive operation scheme using at any time the optimal network configuration of the current demand matrix consumes 17.0 MWh per working day and 15.7 MWh per weekend day; this corresponds to savings of 27.7 % and 33.0 %, respectively.

Figure 11 shows how the power consumption of the network depends on the total traffic. Power consumption in the network scales almost linearly with the total amount of traffic. The slightly disproportionate rise of power when traffic increases from around 4 to 5 Tbps is caused by the need of installing fabric card shelves when more than one line card shelf is necessary: For example, on a working day at 7:00 traffic demand amounts to 4.1 Tbps, and 6 IP routers had active fabric card shelves; this number more than doubles to 14 half an hour later at 7:30 while total traffic demand increases to only 5.3 Tbps.

In contrast to the core network, the open-ring structures of the aggregation network do not allow for rerouting. Still, parallel optical links can be switched off while for each logical connection at least one optical link has to be kept running to avoid violation of the resiliency

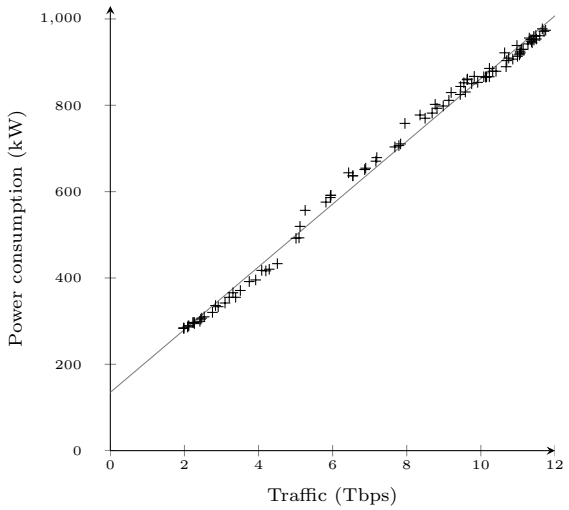

Figure 11: Total power consumption of the network plotted against total amount of traffic (black) with linear regression line (gray).

rules. Figure 12 shows the number of running parallel optical links per logical link in the open-rings of the aggregation network as well as in the core network over a working day. In the core network this number reduces considerably at low traffic times while in the aggregation sections it varies much less; this is a hint that the flexibility in the core network induces a higher (relative) energy-saving potential.

3.2 Energy-optimized switching strategies

Half-hourly reconfiguration of the whole network is probably not desired from an operational point of view. One reason is that newly established connections or light paths are not available immediately after the necessary devices have been switched on, and that changing the traffic routing to use the new connections may lead to instabilities in the network. Another issue is that actual traffic demand might deviate from the nominal values as given in the last section; a thorough way to deal with this would be to apply advanced robust optimization techniques (see, for instance, [4, 2]), which is, however, beyond the scope of this article.

Here we consider several, more simple strategies for adapting the optimization procedure to avoid instabilities in the network:

- Ensure that reconfiguration takes place not too often per day.
- Impose a certain minimal running time of every network configuration.
- Provide for some time in which devices can power up before the new configuration can be used to route traffic.
- Safeguard against temporal deviations from the expected traffic pattern by demand over-provisioning.

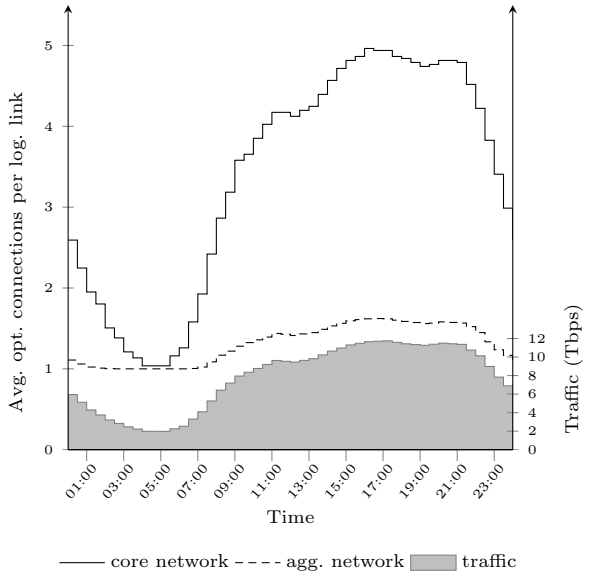

Figure 12: Average number of parallel connections per logical link in the core (solid) and aggregation (dashed) networks for the traffic (gray) over one working day.

Limited number of reconfigurations

Exactly k configurations, each an optimal static configuration for a designated time interval, are allowed per day. Since the traffic data comes in half-hourly time steps, we have $1 \leq k \leq 48$, where $k = 1$ corresponds to the static network where no reconfiguration takes place at all. Clearly, the larger the parameter k, the less energy is consumed over a day: From a given set of k configurations, valid (yet not necessarily optimal) sets of $k + 1$ configurations can be derived by splitting the designated time interval of each configuration in two time intervals and using an optimal configuration in each of the two intervals instead of the original one; each of the resulting sets has a total energy consumption at most as high as the initial set of configurations. Thus, with $k = 48$ the maximal possible energy-savings are achieved.

The problem of finding k energetically favorable time intervals can be reformulated as a shortest path problem in a certain graph under the assumption that power consumption of an optimal network configuration depends affinely linear on the routed traffic (cf. Figure 11). Figure 10 also shows the optimal power levels in a 6-configuration scenario. The possible energy-savings in the backbone for different k on a working day can be read off from columns 2 and 3 in Table 2. It turns out that already with relatively few reconfigurations, say 4, the lion's share of the possible reduction in energy consumption can be realized.

Minimal running time of network configurations

Additionally, we can introduce waiting times t_{wait} between subsequent reconfigurations, that is, each configuration is required to run for a time t_{wait} at the least. With our time discretization model, a waiting time of $t_{\text{wait}} = 1/2\,\text{h}$ imposes no further restriction. Obviously, postulating a waiting time of t_{wait} induces a maximum number of reconfigurations of $k_{\max} = \left\lfloor \frac{24\,\text{h}}{t_{\text{wait}}} \right\rfloor$. Table 2 shows possible energy-savings for different choices of waiting times of 4 h, 6 h, and 8 h. Surprisingly, a waiting time of 6 h does not have too big an impact on energy consumption, compared

Table 2: Energy consumption (MWh) of optimal network configurations for the backbone under different switching regimes (reconfiguration limit k and waiting time t_{wait}) over one working day and energy-savings (%) compared to static configuration.

$t_{\text{wait}} \rightarrow$	½ h		4 h		6 h		8 h	
$k \downarrow$	MWh	%	MWh	%	MWh	%	MWh	%
1	23.4	0.0	23.4	0.0	23.4	0.0	23.4	0.0
2	19.8	15.3	19.8	15.3	19.8	15.3	20.3	13.3
3	19.2	17.9	19.2	17.9	19.2	17.9	20.2	13.6
4	18.7	20.2	19.1	18.5	19.3	17.5		
5	18.5	21.0	19.1	18.4				
6	18.2	22.2	19.1	18.3				
⋮								
48	17.0	27.7						

to other scenarios with maximally 4 reconfigurations allowed per day. An explanation is that the traffic curve (and therefore the power consumption) shows a relatively distinct valley of roughly 6 hours' length and, consequently, the result for 8 h waiting time is quite different.

Another effect of waiting times is that with increasing number of scenarios the total energy consumption not necessarily decreases any more: Splitting the time interval designated to a configuration might lead to an infeasible time interval which violates the waiting time, hence an optimal set of $k + 1$ configurations may have a higher energy consumption than the best sets of k configurations.

Switching times of hardware

Under the assumption that hardware has to be switched on a certain time interval t_{up} before it can be used and can only be switched off after a time interval t_{down} after it is not used any more, it consumes the same amount of power t_{up} before and t_{down} after each reconfiguration. Taking this into account, the total energy consumption in the previously computed scenarios increases slightly, see Table 3.

Table 3: Energy consumption (MWh) of previously considered switching regimes with switching times $t_{\text{up}} = 1/2$ h and $t_{\text{down}} = 1/2$ h.

$t_{\text{wait}} \rightarrow$	½ h		4 h		6 h		8 h	
$k \downarrow$	MWh	%	MWh	%	MWh	%	MWh	%
1	23.4	0.0	23.4	0.0	23.4	0.0	23.4	0.0
2	20.3	13.3	20.3	13.3	20.3	13.3	20.6	12.0
3	19.7	16.0	19.7	16.0	19.7	16.0	20.5	12.4
4	19.1	18.3	19.5	16.7	19.8	15.7		
5	19.0	19.0	19.6	16.5				
6	18.7	20.2	19.6	16.2				

Modification of demand by temporal overprovisioning

To safeguard against uncertainty of traffic volume and time trend, robust optimization techniques can be applied to the integer program above. A simple possibility that avoids making changes to the problem formulation and thereby increasing the problem complexity is to extend the dynamic traffic matrices obtained as described in Section 2 to a set of *temporally robust traffic matrices*. Network configurations computed with these traffic patterns as input are then robust against all dominated traffic situations. Note that the traffic matrices from Section 2 already provide an extra buffer of 20 % of the traffic at each time point. This may not be sufficient to capture temporal deviations from the assumed traffic pattern, though; an expected rise of traffic may occur earlier on some days, or an expected decrease may be delayed longer than anticipated.

We use the following rule to obtain temporal robust traffic matrices. If demand increases between two time steps $t-1$ and t then the demand at t is shifted by time t_{shift} to the previous time step; accordingly for decreasing demand. More precisely, let d_k^t be the traffic demand of some type between a node pair $k = \{a, b\}$ at time step t. Then, $\tilde{d}_k^t = \max\{d_k^{t-t_{\text{shift}}}, \ldots, d_k^{t+t_{\text{shift}}}\}$ is the temporally robust demand between these nodes at time t. The resulting temporally robust traffic matrix for time step t is then obtained again by summing over the matrices $\tilde{d}^t$ for each traffic type. Figure 13 shows the resulting total amount of traffic over a working day for a shift of one hour.

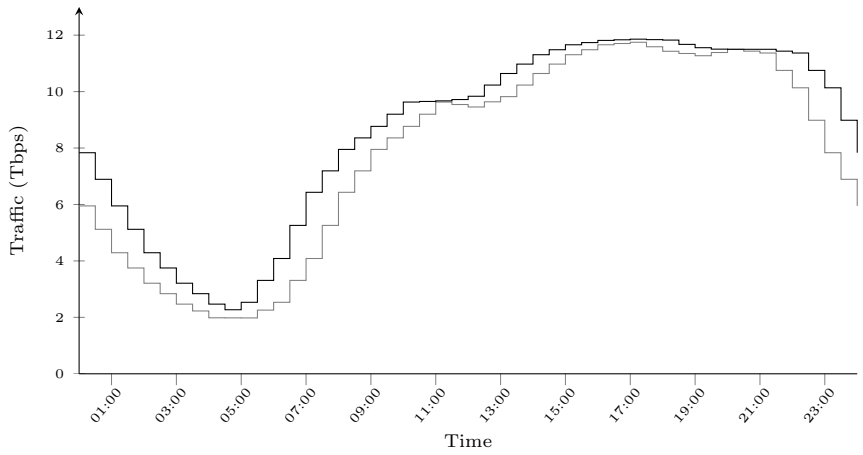

Figure 13: Total traffic curve (black) for a working day from robust traffic matrices obtained by shifting original traffic demands (gray) by $t_{\text{shift}} = 1\,\text{h}$.

The resulting total energy consumption for up to 6 allowed switchings per day with no constraints on waiting time is shown in Figure 14 together with the results for the original demand from Table 2. As can be seen, the energy-saving potential reduces by about 4 – 5 percentage points compared to the results from the original demand. Again, allowing more than 4 reconfigurations does not increase the energy-saving potential by much. For the static scenario, total energy consumption is slightly higher than in the original case, due to the fact that the modified traffic curves for the individual types of traffic add up to a slightly higher value at the peak.

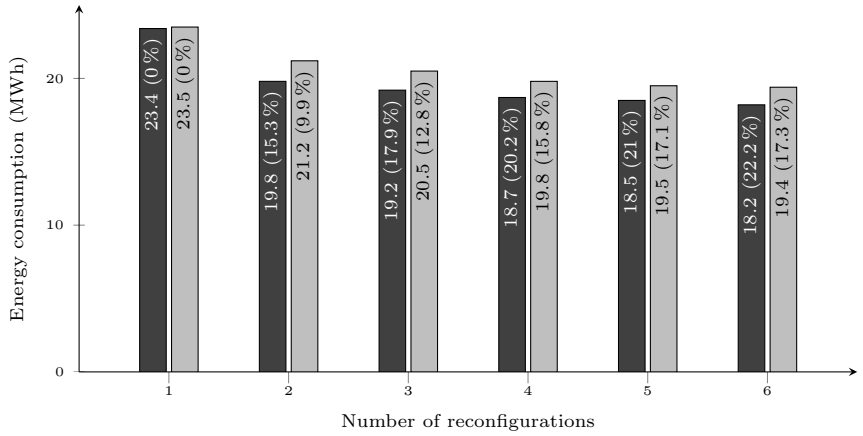

Figure 14: Energy consumption and relative savings resulting from the original demand (dark gray) and the demand shifted by 1 hour (light gray), respectively, with up to 6 reconfigurations per day.

4 Conclusion

We presented a network topology and traffic model taking into account current and future technical and operational requirements. The topology reflects the way networks are designed today to be still sustainable in the future and includes a core network as well as aggregation network sections. The detailed traffic model exhibits traffic values that are realistic for a contemporary network and includes the temporal variation over a working and a week end day.

The network topology and traffic model can be used for various studies. As an example, we presented a model-based optimization approach to assess the potential for energy-savings in a nationwide core network, taking different considerations into account. In a similar way, further network design problems can be dealt with, such as minimizing network instability or classically providing cost-efficient networks provisioned for future traffic requirements.

The detailed network topology together with the traffic data can be requested from the authors (email: `andreas.betker@telekom.de`).

Acknowledgment

This research has been supported by the German Federal Ministry for Economic Affairs and Energy (BMWi) within the IT2Green programme (`www.it2green.de`) and the German Federal Ministry of Education and Research (BMBF) within the EUREKA project SASER.

References

[1] B. Addis, A. Capone, G. Carello, L.G. Gianoli, and B. Sanso. Energy management through optimized routing and device powering for greener communication networks. *IEEE/ACM Transactions on Networking*, 22(1):313–325, February 2014.

[2] Thomas Bauschert, Christina Büsing, Fabio D'Andreagiovanni, Arie M. C. A. Koster, Manuel Kutschka, and Uwe Steglich. Network planning under demand uncertainty with robust optimization. *IEEE Communications Magazine*, to appear.

[3] Berlin Commercial Internet Exchange. Statistics, 2013. http://www.bcix.de/bcix/statistik/.

[4] Dimitris Bertsimas, David B. Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

[5] A. Betker and C. Gerlach. Traffic demand model for national optical networks based on the population distribution. In *ITG-Fachtagung Photonische Netze*, pages 71–76, 2003.

[6] A. Betker, C. Gerlach, R. Hülsermann, M. Jäger, M. Barry, S. Bodamer, J. Späth, C. Gauger, and M. Köhn. Reference transport network scenarios. Technical report, MultiTeraNet, 2003. Available online: http://www.ikr.uni-stuttgart.de/IKRSimLib/Usage/Referenz_Netze_v14_full.pdf.

[7] Andreas Betker, Dirk Kosiankowski, Christoph Lange, Frank Pfeuffer, Christian Raack, and Axel Werner. Energy efficiency in extensive multilayer core and regional networks with protection. Technical Report 12-45, ZIB, Takustr.7, 14195 Berlin, 2012.

[8] Andreas Betker, Dirk Kosiankowski, Christoph Lange, Frank Pfeuffer, Christian Raack, and Axel Werner. Energy efficiency in extensive IP-over-WDM networks with protection. In Stefan Helber, Michael Breitner, Daniel Rösch, Cornelia Schön, Johann-Matthias Graf von der Schulenburg, Philipp Sibbertsen, Marc Steinbach, Stefan Weber, and Anja Wolter, editors, *Operations Research Proceedings 2012*, pages 93–99. Springer, 2014.

[9] BITKOM. Markt & Statistik. http://www.bitkom.org/de/markt_statistik/64002.aspx [last accessed on 2014-05-23].

[10] Cisco. Visual Networking Index: Forecast and Methodology. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html.

[11] A. Dwivedi and R. E. Wagner. Traffic model for USA long distance optical network. In *Optical Fiber Communication Conference (OFC)*, pages 156–158, 2000.

[12] German Commercial Internet Exchange. Traffic statistics, 2013. http://www.de-cix.net/about/statistics/.

[13] W. Van Heddeghem and F. Idzikowski. Equipment power consumption in optical multilayer networks – source data. Technical report, Ghent Univ., 2012. http://powerlib.intec.ugent.be.

[14] W. Van Heddeghem, F. Idzikowski, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester. Power consumption modeling in optical multilayer networks. *Photonic Network Communications*, 2012. published online: 26 Jan 2012.

[15] R. Huelsermann, M. Gunkel, C. Meusburger, and D. A. Schupke. Cost modeling and evaluation of capital expenditures in optical multilayer networks. *Journal of Optical Networking*, 7(9):814–833, 2008.

[16] F. Idzikowski, S. Orlowski, Ch. Raack, H. Woesner, and A. Wolisz. Dynamic routing at different layers in IP-over-WDM networks – Maximizing energy savings. *Optical Switching and Networking, Special Issue on Green Communications*, 8(3):181–200, 2011.

[17] C. Lange, D. Kosiankowski, A. Betker, H. Simon, N. Bayer, D. v. Hugo, H. Lehmann, and A. Gladisch. Energy efficiency of load-adaptively operated telecommunication networks. *Journal of Lightwave Technology*, 32(4):571–590, 2014. invited paper.

[18] Ulrich Menne, Christian Raack, Roland Wessäly, and Daniel Kharitonov. Optimal degree of optical circuit switching in IP-over-WDM networks. In *Proceedings of the 16th International Conference on Optical Network Design and Modeling (ONDM)*, pages 1–6, 2012.

[19] S. Orlowski, R. Wessäly, M. Pióro, and A. Tomaszewski. Sndlib 1.0—survivable network design library. *Networks*, 55(3):276–286, May 2010. `http://sndlib.zib.de`.

[20] Giuseppe Rizzelli, Annalisa Morea, Massimo Tornatore, and Olivier Rival. Energy efficient traffic-aware design of on–off multi-layer translucent optical networks. *Computer Networks*, 56(10):2443 – 2455, 2012.

[21] Statistisches Bundesamt. German statistical yearbook 2013, 2013. `https://www.destatis.de/DE/Publikationen/StatistischesJahrbuch/StatistischesJahrbuch.html` [last accessed on 2014-05-23].

[22] STRONGEST. Scalable, tunable and resilient optical networks guaranteeing extremely-high speed transport, 2010. `http://www.ict-strongest.eu/upload-files/newsletters-1/the-strongest-news-no1-5/jlist-frontend-header-finish-title` [last accessed on 2014-05-23].

[23] 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, January 2006.