Dynamic Resource Operation and Power Model for IP-over-WSON Networks

Uwe Bauknecht, Frank Feller

Dynamic Resource Operation and Power Model for IP-over-WSON Networks

Uwe Bauknecht and Frank Feller

Institute of Communication Networks and Computer Engineering (IKR)
Universität Stuttgart, Pfaffenwaldring 47, 70569 Stuttgart, Germany
{uwe.bauknecht,frank.feller}@ikr.uni-stuttgart.de

Abstract. The power consumption of core networks is bound to grow considerably due to increasing traffic volumes. Network reconfiguration adapting resources to the load is a promising countermeasure. However, the benefit of this approach is hard to evaluate realistically since current network equipment does not support dynamic resource adaptation and power-saving features. In this paper, we derive a dynamic resource operation and power model for IP-over-WSON network devices based on static power consumption data from vendors and reasonable assumptions on the achievable scaling behavior. Our model allows to express the dynamic energy consumption as a function of active optical interfaces, line cards, chassis, and the amount of switched IP traffic. We finally apply the model in the evaluation of two network reconfiguration schemes.

Keywords: Dynamic Energy Model, Resource Model, IP-over-WDM Network, Multi-layer Network, Energy Efficiency

1 Introduction

The continuous exponential growth of the traffic volume is driving an increase in the power consumption of communication networks, which has gained importance due to economic, environmental, and regulatory considerations. While the access segment has traditionally accounted for the largest share of the energy consumed in communication networks, core networks are bound to become predominant due to new energy-efficient broadband access technologies like FTTx [1, 2].

Today, core network resources are operated statically whereas the traffic exhibits significant variations over the day, including night periods with as little as 20 % of the daily peak rate [3]. Adapting the amount of active, i. e. energy-consuming, resources to the traffic load could therefore significantly reduce power consumption. Since the paradigm of static operation has traditionally guaranteed the reliability of transport networks, research needs to thoroughly investigate both the benefits and the consequences of dynamic resource operation.

A first challenge is to determine the energy savings achievable by this approach. Current core network equipment neither allows the deactivation of components nor features state-of-the-art power-saving techniques (such as frequency

scaling found in desktop processors). Hence, its power consumption hardly depends on the traffic load [4, 5]. For the static operation of network equipment, researchers have derived power models from measurements [4] and product specifications or white papers [6]. These models are either intended as a reference for future research [6] or applied in power-efficient network planning studies [7].

Core network reconfiguration studies base on different power adaptation assumptions. Mostly, components (e. g. optical cross-connects, amplifiers, regenerators [8, 9]; router interfaces, line cards [10, 11]) or whole links and nodes [10] are switched off. Alternatively, the energy consumption is scaled with active capacity [12], link bit rate [13], or traffic load [14]. Restrepo *et al.* [15] propose load-dependent energy profiles justified by different power-saving techniques for components and also apply them to nodes. Lange *et al.* [16] derive a dynamic power model from a generic network device structure and validate it by measurements of one piece of power-saving enabled equipment. To our knowledge, however, there is no coherent reference model for the dynamic power consumption of current core network equipment assuming state-of-the-art power-scaling techniques. Such a model is highly desirable for network reconfiguration studies.

In this paper, we derive a dynamic resource operation and power model for Internet protocol (IP) over wavelength switched optical network (WSON) networks. We start by decomposing network nodes into their relevant hardware components, for which we obtain static power consumption values from reference models and publicly available manufacturer data. Based on their structure and functions, we identify applicable activation/deactivation patterns and power-scaling techniques and derive the operation state-dependent power consumption. We then aggregate the component behavior into a power model suitable for network-level studies. It features power values for optical circuits and electrically switched traffic. In addition, it takes the resource hierarchy of network nodes into consideration and describes their static power consumption. We finally illustrate the application of our model using reconfiguration schemes from [11, 17].

2 Network Equipment Structure

The two layers of IP-over-WSON networks are reflected in the structure of the network nodes. The nodes typically consist of a router processing IP packets for the upper layer and a wavelength-selective optical circuit switch in the lower layer. Optical circuits interconnect arbitrary IP routers.

2.1 IP Routers

Large core routers like the Cisco CRS-3, Alcatel-Lucent (ALU) 7950 XRS, or Juniper T4000 have a modular structure as depicted in Fig. 1. Their basic element is the line card chassis (LCC), which features a passive backplane interconnecting card slots. The chassis is equipped with a number of elements to provide basic functionality. This comprises power supply and cooling units realized in a modular way for failure tolerance and step-wise upgrades. Power supply and cooling

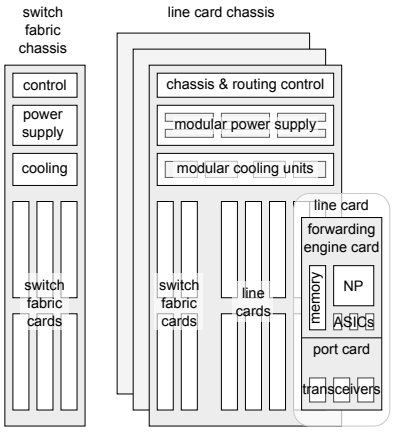

Fig. 1. IP layer device structure

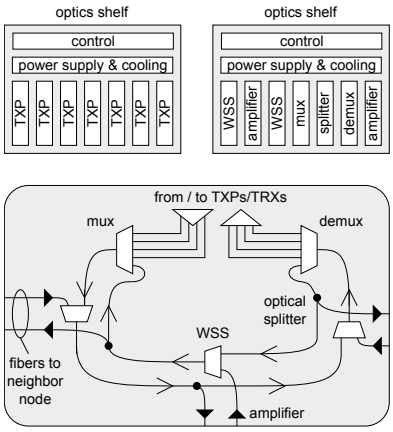

Fig. 2. WSON node structure

are managed by the chassis control (CC) among other tasks. The routing engine (RE) exchanges routing protocol messages with other network nodes and builds a routing table. Essentially, the CC and RE modules are small general-purpose server systems. In the case of Cisco, one such system fulfills both functions.

Line cards (LC) perform the central role of the router: they terminate optical circuits and switch the traffic on the IP packet level. One can subdivide them according to these functions into port cards (PC) and forwarding engine cards (FE). The latter feature network processors (NP), further application specific integrated circuits (ASIC), and memory to process and store packets. NPs generally have highly parallel architectures[1]. Typically, one part of the memory is used for packet buffering, whereas a copy of the routing table occupies the rest.

PCs provide the connection between optical circuits and the electrical interfaces of FEs. One PC may terminate several circuits. Traditionally, PCs feature ports for pluggable optical transceiver (transmitter and receiver) modules such as SFP[2], XFP[3], or CFP[4]. These transceivers (TRX) are so-called short-reach transceivers (SR-TRX) of a limited optical reach (100 m to 100 km); they connect the LCs to transponders (TXP) converting the signal to a given wavelength for long-haul transmission. Alternatively, PCs can directly feature *colored* interfaces with long-reach transceivers (LR-TRX). This avoids the overhead of signal conversion, but currently restricts the number of interfaces per LC.

Switch fabric cards (SFC) comprising ASICs and buffer memory interconnect LCs over the passive backplane. They are essential for the forwarding of packets between different LCs. The total interconnection capacity between any pair of LCs is distributed over several SFCs, allowing graceful degradation in case of failures. It is likewise possible to interconnect several LCCs using a dedicated

[1] Cf. Cisco Quantum Flow Array [18], ALU FP3 [19], EZchip NPS-400 [20]
[2] Small form-factor pluggable, capable of 1 Gbps
[3] 10 Gigabit small form-factor pluggable
[4] C form-factor pluggable, capable of 100 Gbps

switch fabric card chassis (FCC), which features further SFCs, in order to create a logical router of higher capacity.

2.2 WSON Devices

Functionally, devices of the WSON layer are subdividable into generation/termination of optical signals and their switching and transmission. The former task is performed by TXPs and colored LC interfaces. On the long-reach side, these devices feature lasers for transmission (and as local oscillators in case of coherent detection), photodiodes, digital-to-analog converters (DAC), analog-to-digital converters (ADC), and digital signal processing (DSP) ASICs. In addition to forward error correction (FEC), the DSP performs the complex task of compensating for optical impairments in the case of high bit-rate channels.

In WSON, optical switching nodes transfer wavelength-division multiplex (WDM) signals between fibers connecting to neighbor nodes and from/to local TXPs/TRXs. There exists a variety of configurations differing in functionality and complexity. We assume the setting depicted in the lower part of Fig. 2. It is *colorless* and *directionless*, meaning that any incoming WDM signal is switchable to any output fiber (as long as there is only one fiber pair to each neighbor node). Each wavelength may only be used once in the node. Key components are optical splitters relaying channels to neighbor nodes and wavelength selective switches (WSS), which are able to select any wavelength from either the local ring or the incoming fiber. In addition, WSSs act as multiplexer (mux) and demultiplexer (demux) to insert and drop locally terminated channels.

Optical (i. e. analog) signal amplification is not only needed at the input and output of the optical node, but also along the fibers. Typically, an optical line amplifier (OLA) is placed every 80 km. Introducing noise, this in turn limits the distance without electrical (i. e. digital) signal regeneration (essentially by two back-to-back TXPs) to between 800 km and 4000 km.

Like in the IP layer, the optical components present in nodes are organized in shelves providing power supply, cooling, and control (cf. Fig. 2, top). Such shelf systems are e. g. ALU 1830 Photonic Service Switch, ADVA FSP 3000 or Cisco ONS 15454 Multiservice Transport Platform.

3 Static Power Consumption Values

We derive our power model from static (maximum) power values for Cisco's CRS-3, ONS 15454 MSTP, and their respective components, as well as a stand-alone EDFA by Finisar. We mainly refer to Cisco equipment since to our knowledge, Cisco is the only vendor to publicly provide power values of individual components for both core router and WDM equipment. This also makes our work comparable to that of other researchers who use the same data for similar reasons [6, 7, 35]. Table 1 lists the static power values (along with variables referring to them in Sect. 4.2). They are largely identical to the ones in the Powerlib[5].

[5] Cf. `http://powerlib.intec.ugent.be` and [6]

Table 1. Static power values per component. Symbols are explained in section 4.

Component type	Power Consumption	Number installable	Source
FCC control per chassis	1068 W $/\eta := P_{\mathrm{SCC}}$	1 for 9 LCC	[21]
Switch Fabric Card for FCC	229 W $/\eta := P_{\mathrm{S2SFC}}$	8 for 3 LCC	[21]
Optical Interface Module	166 W $/\eta := P_{\mathrm{OIM}}$	8 for 3 LCC	[21]
Chassis Control	275 W $/\eta := P_{\mathrm{CC}}$	2 per LCC	[22]
Fan Tray	344 W $/\eta := P_{\mathrm{CU}}$	2 per LCC	[6]
Switch Fabric Card for LCC	206 W $/\eta := P_{\mathrm{SFC}}$	8 per LCC	[6]
140G Forwarding Engine	446 W $/\eta := P_{\mathrm{FE}}$	1 to 16 per LCC	[6]
1x 100G LR-TRX	180 W $/\eta \in \boldsymbol{P}_{\mathrm{TRX}}$	1 per FE	[23]
2x 40G LR-TRX	150 W $/(2 \cdot \eta) \in \boldsymbol{P}_{\mathrm{TRX}}$	1 per FE	[6]
4x 10G LR-TRX	150 W $/(4 \cdot \eta) \in \boldsymbol{P}_{\mathrm{TRX}}$	1 per FE	[6]
1x 100G SR-port card	150 W $/\eta \in \boldsymbol{P}_{\mathrm{PC}}$	1 per FE	[6]
2x 40G SR-port card	185 W $/\eta \in \boldsymbol{P}_{\mathrm{PC}}$	1 per FE	[24]
4x 40G SR-port card	185 W $/\eta \in \boldsymbol{P}_{\mathrm{PC}}$	1 per FE	[25]
14x 10G SR-port card	150 W $/\eta \in \boldsymbol{P}_{\mathrm{PC}}$	1 per FE	[6]
20x 10G SR-port card	150 W $/\eta \in \boldsymbol{P}_{\mathrm{PC}}$	1 per FE	[6]
1x 100G SR-TRX (CFP)	12 W $/\eta \in \boldsymbol{P}_{\mathrm{TRX}}$	1 per 100G PC	[26]
1x 100G SR-TRX (CXP)	6 W $/\eta \in \boldsymbol{P}_{\mathrm{TRX}}$	1 per 100G TXP	[27]
1x 40G SR-TRX	8 W $/\eta \in \boldsymbol{P}_{\mathrm{TRX}}$	1 to 4 per 40G PC	[28]
1x 10G SR-TRX	3.5 W $/\eta \in \boldsymbol{P}_{\mathrm{TRX}}$	1 to 14 per 10G PC	[27]
1x 100G TXP (excl. SR-TRX)	133 W $/\eta \in \boldsymbol{P}_{\mathrm{TXP}}$	6 per 100G-Shelf	[29]
1x 40G TXP (incl. SR-TRX)	129 W $/\eta \in \boldsymbol{P}_{\mathrm{TXP}}$	6 per shelf	[30]
1x 10G TXP (excl. SR-TRX)	50 W $/\eta \in \boldsymbol{P}_{\mathrm{TXP}}$	12 per shelf	[6]
OLA standalone	50 W $/\eta := P_{\mathrm{OLA}}$	1 every 80 km	[31]
OLA card	40 W $/\eta := P_{\mathrm{Amp}}$	2 on every edge	[32]
WSS card	20 W $/\eta := P_{\mathrm{WSS}}$	1 on every edge	[33],[6]
12-Port Optics Shelf	260 W $/\eta := P_{\mathrm{OS}}$	-	[34],[33]
6-Port Optics Shelf (100G)	284 W	-	[34]

In the CRS-3, the components are combined as follows: One LCC houses one to eight power supply modules, which can reach an efficiency of about 95 % [21]. In addition, there are two fan trays and two CC cards (also acting as RE). The switch fabric is realized by an output-buffered Beneš network distributed on 8 parallel SFCs. A multi-chassis configuration requires an FCC with a set of control modules and eight special SFCs and optical interface modules (OIM) per group of three LCCs [7]. The LCCs then use a different type SFCs with a similar power consumption. We use the most powerful type of FE (or *Modular Services Card*) capable of 140 Gbps, which allows for PCs with a maximum of 1x 100G, 3x 40G or 14x 10G ports at full line rate. On the optical layer, we consider TXPs matching these performance values. While the 40G TXP has an integrated SR-interface, the 100G TXP and 10G TXP need additional SR-TRXs.

For the ONS 15454 MSTP version with 100G-capable backplane, Cisco lists a power consumption of 284 W for shelf, cooling, and controller card [34]. We estimate the power consumption of the larger 12-port version, for which no such reference is available, to 260 W, since both cooling and CC are slightly less power-hungry for this version (although exact numbers vary between references[6]). Contrary to the specification, we assume that the 12-port version can hold the 100G TXP as long as the backplane is not used. Unlike the stand-alone OLA, the OLA modules for the shelves are unidirectional. Consequently, two of them are deployed for each connection to another router. The WSS is Cisco's 80-channel WXC, which additionally comprises a passive beam splitter.

[6] Cf. tables A-1, A-4 and A-5 from [34] and table A-1 from [33]

4 Dynamic Resource Operation and Power Model

We derive a dynamic model for the momentary power consumption of a hypothetical core network node with extended power saving features. For the application in network-level studies, we limit the scope of power scaling to adapting processing capacity to the amount of packet-switched traffic and modifying optical circuits (along with the hierarchically required resources like LCs and LCCs). The according scaling mechanisms operate at different time scales: To follow fast fluctuations of the packet rate, processing capacity needs adaptation in the order of microseconds. In contrast, establishing or tearing down an optical circuit takes in the order of several minutes due to transient effects in OLAs. We assume that this latter time scale enables the (de)activation of any electronic component.

In the following, we discuss the power scaling possibilities of the node's components. We then describe the resulting model.

4.1 Dynamic Operation and Power Scaling Assumptions

Power Supply: We expect the CRS power supply modules to behave like the most efficient Cisco models, which reach an efficiency of at least $\eta = 90\,\%$ in a load range of 25 % to 90 % [36]. Since chassis are operable with a varying number of such modules, we assume that we can switch modules off to stay in the 90 % efficiency region. We accordingly derive the gross power consumption by increasing all power values by a factor of $1/\eta \approx 1.11$.

Cooling: The affinity laws relate the power consumption P_i of fans at operation point i to their rotational speed N_i by $P_1/P_2 = (N_1/N_2)^3$. The normal operation range of CRS fans is 3300 RPM to 5150 RPM, with the maximum rated at 6700 RPM [37]. We obtain a net minimum power consumption of $(3300/6700)^3 \approx 12\,\%$ of the total fan tray power. Accounting for driver overhead, we estimate the static fraction at 20 % and let the remainder scale linearly with the number of active LCs, since these produce the largest share of heat.

Chassis Controller Cards: Being a small general-purpose server system, the CC can readily benefit from power saving features like dynamic voltage and frequency scaling (DVFS) or power-efficient memory [38]. We do however not model the control workload and consequently consider that the CC consume constant (maximum) power.

Switch Fabric: The interconnection structure prohibits the scaling of SFCs along with LCs, and the signaling time overhead for SFC (de)activation impedes an adaptation to the switched IP traffic. We therefore assume all SFCs to be active when their LCC is so, and we disregard power scaling options of ASICs and memory *on* SFCs. For multi-chassis routers, we allow active SFCs and OIMs in the FCC to scale with LCCs in accordance with possible static configurations (i. e. we switch blocks of eight SFCs and OIMs per group of three LCCs).

Forwarding Engine Cards: ASICs and NPs account for 48 % of the power budget of a LC, memory for 19 %; the remaining 33 % is spent on power conversion, control and auxiliary logic [39]. We assume the latter part to be static. The same applies to 9 % (out of the 19 %) of memory power consumption for the forwarding information base [40]. The remainder (10 % of the LC consumption) is for buffer memory, which is presumably dimensioned for the 140 Gbps capacity of the FE following the bandwidth-delay product (BDP) rule. We let the active buffer memory scale according to the BDP with the capacity of the active circuits terminated by the LC. Neglecting the residual power consumption in deep sleep state and the discrete nature of switchable memory units [40], we assume the buffer's power consumption to scale linearly with the active port capacity.

Recent NP designs support power saving by switching off unused components, e. g. cores [19, 41, 42]. In theory, deactivating and applying DVFS to NP cores alone can save more than 60 % in low traffic situations [43]. We therefore assume that 70 % of the power consumption of ASICs and NPs dynamically scales with processed IP traffic, while 30 % is static. In sum, we assign $33\,\% + 9\,\% + (30\,\% \cdot 48\,\%) = 56.4\,\%$ statically to an active LC; we let 10 % scale with the active port capacity and $70\,\% \cdot 48\,\% = 33.6\,\%$ with processed traffic.

Port Cards, Transceivers, Transponders: Like the latest hardware generation [16, 41], we allow the dynamic deactivation of unused LC ports. We assume that a (multi-port) PC is active along with the FE as long as one of its ports is so. TRXs and TXPs are switched on and off with the respective circuits. We disregard more fine-grained power scaling proposals for TXP ASICs [44], but we do assume that the transmit and receive parts of a TXP may be activated separately[7]. In case of a PC with multiple colored interfaces, we distribute the total power consumption of the PC to the ports and let it scale with the circuits.

Optical Node: We consider the power consumption for cooling, power supply, and control of an optics shelf as static, since it is much less than the consumption of the respective LCC modules. The same applies to OLAs and WSSs. We do however allow the deactivation of TXP-hosting shelves when all TXPs are switched off.

4.2 Resulting Power Model

For the mathematical model, we use the following conventions: Capital letters without indexes represent equipment-specific constants. C_α denotes capacity in Gbps, P_α maximum power consumption and N_α the maximum installable quantity of component α. The set of all possible values for P_α is denoted by $\boldsymbol{P}_\alpha$. Indexed variables denote a specific component in a specific node. Variables of the type n_β represent dimensioning parameters describing the configuration of one node. Small letters indicate model parameters characterizing the dynamic configuration and load situation.

[7] The power distribution between the parts is of minor relevance since one of each is needed for one circuit.

A router may consist of a maximum of $N_{\mathrm{LCC}} = 9$ LCCs; the actual number is $n_{\mathrm{LCC}} \in \{1..N_{\mathrm{LCC}}\}$. If the router has more than one LCC, a FCC is needed. In this case the first factor of (1) is nonzero, and so is $P_{\mathrm{FCC,stat}}$.

A LCC is always equipped with $N_{\mathrm{SFC}} = 8$ SFCs, $N_{\mathrm{CC}} = 2$ CC cards and $N_{\mathrm{CU}} = 2$ cooling units. We consider 20 % of the cooling power consumption as static, resulting in the total static LCC consumption in (2). The remaining 80 % of the cooling units gives the dynamic power consumption of LCC i with x_i active out of $n_{\mathrm{LC},i} \in \{1..N_{\mathrm{LC}}\}$ installed LCs according to (3). One LCC houses at most $N_{\mathrm{LC}} = 16$ LCs. An active LC statically consumes 56.4% of P_{FE} and the total of the respective port card out of $\boldsymbol{P}_{\mathrm{PC}}$. The power cost of the j-th LC in LCC i is thus given by (4). The power consumption per active port has two components: The packet buffers of the FE and the respective TRX $P_{\mathrm{TRX},ij} \in \boldsymbol{P}_{\mathrm{TRX}}$. For colored LC interfaces, $P_{\mathrm{TRX},ij}$ is the consumption of the PC over the number of ports and $P_{\mathrm{PC},ij} = 0$. The buffer's consumption (10 % of P_{FE}) is scaled to the capacity of one port $C_{\mathrm{P},ij}$ ($C_{\mathrm{FE}} = 140$ Gbps is the total capacity of the FE). The dynamic LC power share is obtained by multiplying these contributions by the number of active ports y_{ij}, (5). Equation (6) represents the traffic-dependent power consumption of a LC. The variable z_{ij} indicates the traffic demand in Gbps at LC j in LCC i.

In the optical layer, we assume a ring configuration according to Fig. 2. Each of the n_{d} links to a neighbor node requires one WSS module and two amplifiers. The router is connected through two additional WSS modules. WSS modules, amplifiers and TXPs each use one of 12 available slots in an optics shelf. This adds up to the static power of optical components $P_{\mathrm{Opt,stat}}$ in (7), where n_λ is the maximum number of wavelengths needed on one link, $N_\lambda = 80$ the maximum number of wavelengths per fiber and P_{WSS}, P_{Amp} and P_{OS} are the power consumptions of a WSS module, OLA module and optics shelf respectively. n_{TXP} is the number of installed TXPs.

$$P_{\mathrm{FCC,stat}} = \left\lceil \frac{n_{\mathrm{LCC}} - 1}{N_{\mathrm{LCC}}} \right\rceil \left(\left\lceil \frac{n_{\mathrm{LCC}}}{3} \right\rceil \cdot 8(P_{\mathrm{OIM}} + P_{\mathrm{S2SFC}}) + P_{\mathrm{SCC}} \right) \quad . \tag{1}$$

$$P_{\mathrm{LCC,stat}} = N_{\mathrm{SFC}} \cdot P_{\mathrm{SFC}} + N_{\mathrm{CC}} \cdot P_{\mathrm{CC}} + N_{\mathrm{CU}} \cdot P_{\mathrm{CU}} \cdot 20\% \quad . \tag{2}$$

$$P_{\mathrm{LCC,dyn},i}(x_i) = \frac{x_i}{N_{\mathrm{LC}}} \cdot N_{\mathrm{CU}} \cdot P_{\mathrm{CU}} \cdot 80\% \quad . \tag{3}$$

$$P_{\mathrm{LC,stat},ij} = P_{\mathrm{FE}} \cdot 56.4\% + P_{\mathrm{PC},ij} \quad . \tag{4}$$

$$P_{\mathrm{LC,dyn},ij}(y_{ij}) = \left(\frac{P_{\mathrm{FE}} \cdot 10\%}{C_{\mathrm{FE}}} \cdot C_{\mathrm{P},ij} + P_{\mathrm{TRX},ij} \right) \cdot y_{ij} \quad . \tag{5}$$

$$P_{\mathrm{LC,traff},ij}(z_{ij}) = \frac{48\% \cdot 70\% \cdot P_{\mathrm{FE}}}{C_{\mathrm{FE}}} \cdot z_{ij} \quad . \tag{6}$$

$$P_{\mathrm{Opt,stat}} = \left\lceil \frac{n_\lambda}{N_\lambda} \right\rceil \cdot \left(n_{\mathrm{d}} \cdot (P_{\mathrm{WSS}} + 2 \cdot P_{\mathrm{Amp}}) + 2 \cdot P_{\mathrm{WSS}} \right) + \left\lceil \left(\left\lceil \frac{n_\lambda}{N_\lambda} \right\rceil \cdot (n_{\mathrm{d}} \cdot 3 + 2) + n_{\mathrm{TXP}} \right) \cdot \frac{1}{12} \right\rceil \cdot P_{\mathrm{OS}} \quad . \tag{7}$$

Equation (8) finally gives the dynamic consumption $P_{\mathrm{Opt,dyn}}$ of the TXPs connected to LC j in LCC i.

$$P_{\mathrm{Opt,dyn}}(y_{ij}) = \sum_{i=1}^{n_{\mathrm{LCC}}} \sum_{j=1}^{n_{\mathrm{LC},i}} P_{\mathrm{TXP},ij} \cdot y_{ij} \ . \tag{8}$$

5 Application Example

We illustrate the application of the dynamic power model by evaluating the power savings achieved by different degrees of network reconfiguration assuming two different node configurations.

5.1 Scenario

Node Configuration. We assume WDM channels of 40 Gbps, which are terminated either by colored LC interfaces (*case A*) or by TXPs in a WDM shelf connected to LCs via short-reach optics (*case B*). While each LC can only feature one interface in case A, we serve up to three 40 Gbps channels with one LC in case B. In both cases, add/drop traffic is handled by 10 Gbps short-reach interfaces on dedicated tributary LCs (with up to 14 such interfaces). Unlike resources on the core network side, all tributary interfaces are constantly active. Table 2 gives the numerical power values. The port comprises TXP (and SR-TRX in case B) as well as the dynamic line card power share. Values for the optical equipment are used in (7) according to node dimensioning.

Network and Traffic. We present results for the 22-node Géant reference network (available from SNDLib [45]) using ten days out of demand traces collected in 2005 [46]. Assuming similar statistical properties in these ten days, we give 95 % confidence intervals for the metrics. To vary the network load, we linearly scale the demand traces and quantify this scaling based on a *peak demand matrix* containing the peak values of each node-to-node demand trace. We scale the traces such that the average of the values in the peak demand matrix ranges between 2 Gbps and 160 Gbps, corresponding to a total peak demand (sum over all values in the matrix) of 924 Gbps to 73.9 Tbps.

Resource Adaptation. We consider three operation schemes and evaluate the power consumption every 15 minutes: *(i)* In a static network configuration (regarding virtual topology and demand routes optimized for the peak demand matrix), we let active resources scale with the load. This corresponds to FUFL in [11]. *(ii)* We reconfigure the virtual topology by periodically applying the *centralized dynamic optical bypassing* (CDOB) scheme according to [17], using the actual power consumption as cost function with the simulated annealing-based solution method. *(iii)* The static operation of all resources dimensioned for the peak demand matrix serves as baseline.

Table 2. Power values (in Watts)

Component	Contribution	Case A	Case B
FCC	static if >1 LCC	1124	1124
SFC in FCC	per 3 LCC	3326	3326
LCC	if $\geq$ 1 LC on	2336	2336
LC	if $\geq$ 1 port on	286	471
Port	per circuit	163	150
Tributary LC	static	436	436
Tributary port	static	3.5	3.5
IP processing	per Gbps	1.07	1.07
Optics shelf	per 12 units	260	260
WSS+2Ampl.	static	100	100
2 WSS	static	40	40

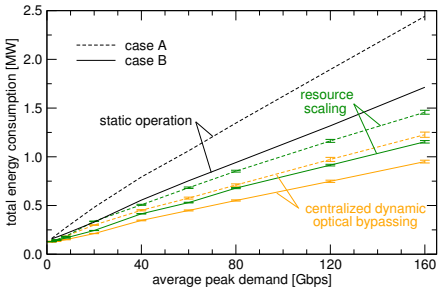

Fig. 3. Network-wide power consumption

5.2 Results

Fig. 3 plots the average power consumption of all devices in the network over the average peak demand for the different configuration cases and adaptation schemes. As one would expect, the power consumption increases with the load in all cases. We further observe that the power consumption is systematically higher for case A compared to the same adaptation scheme in case B: the savings due to the higher port density per LC and per LCC overcompensate the cost for additional short-reach optics.

Energy savings by resource adaptation in an otherwise static network configuration range between 20 % for low average load and 40 % for high load in case A. The increasing benefits are explained by a higher number of parallel circuits allowing deactivation. CDOB likewise requires a certain amount of traffic to be effective and yields 10 additional percentage points of savings over the simple dynamic resource operation at maximum load. In case B, the achievable savings only range between 20 % and 33 % (resp. 44 % for CDOB), albeit at an altogether lower power consumption level. This effect is explained by the resource hierarchy: unlike in case A, LCs may need to remain active when only a fraction of their ports is used.

6 Conclusion

In this paper, we derived a dynamic power model for IP-over-WSON network equipment assuming the presence of state-of-the-art power-saving techniques in current network devices. Based on an in-depth discussion of the power scaling behavior of the components of the devices, we express the dynamic power consumption of a network node as a function of the optical circuits it terminates and the pieces of the hardware hierarchy (line cards, racks) required for this. In addition, a smaller share of power scales with the amount of electrically processed traffic.

Lacking reliable data on power scaling options for optical switching equipment and amplifiers, we assume that the power consumption of these devices

is independent of the load. While the model is open to improvement in this respect, the impact of this limitation is small given the predominance of the power consumption of IP layer equipment and transceivers.

By applying our power model to different network resource adaptation schemes in an example scenario, we found that dynamic resource operation can reduce the total power consumption of the network by 20 % to 50 %. However, generalizing these figures requires much caution since they strongly depend on the assumed resource dimensioning in the static reference case.

Future work could extend the dynamic power model in order to take new devices and trends into account. E. g., the power consumption of rate-adaptive transponders could be modeled. Likewise, different optical node variants could be included. Besides, the model awaits application in research on network reconfiguration.

Acknowledgments. The work presented in this paper was partly funded within the SASER project SSEN by the German Bundesministerium für Bildung und Forschung under contract No. 16BP12202

References

1. Baliga, J., et al.: Energy consumption in optical IP networks. J. Lightwave Technol. **27**(13) (2009)
2. Hinton, K., et al.: Power consumption and energy efficiency in the Internet. IEEE Netw. **25**(2) (2011)
3. De-Cix: Traffic statistics (2013) `http://www.de-cix.net/about/statistics/`.
4. Chabarek, J., et al.: Power awareness in network design and routing. In: Proc. INFOCOM. (2008)
5. Hülsermann, R., et al.: Analysis of the energy efficiency in IP over WDM networks with load-adaptive operation. In: Proc. 12th ITG Symp. on Photonic Netw. (2011)
6. Van Heddeghem, W., et al.: Power consumption modeling in optical multilayer networks. Photonic Netw. Communic. **24** (2012)
7. Wang, L., et al.: Energy efficient design for multi-shelf IP over WDM networks. In: Proc. Computer Communic. Workshops at INFOCOM. (2011)
8. Wu, Y., et al.: Power-aware routing and wavelength assignment in optical networks. In: Proc. ECOC. (2009)
9. Silvestri, A., et al.: Wavelength path optimization in optical transport networks for energy saving. In: Proc. ICTON. (2009)
10. Chiaraviglio, L., et al.: Reducing power consumption in backbone networks. In: Proc. ICC. (2009)
11. Idzikowski, F., et al.: Dynamic routing at different layers in IP-over-WDM networks – maximizing energy savings. Optical Switching and Netw. **8**(3) (2011)
12. Bathula, B.G., et al.: Energy efficient architectures for optical networks. In: Proc. London Commun. Symp. (2009)
13. Vasić, N., et al.: Energy-aware traffic engineering. In: Proc. e-Energy. (2010)
14. Puype, B., et al.: Power reduction techniques in multilayer traffic engineering. In: Proc. ICTON. (2009)

15. Restrepo, J., et al.: Energy profile aware routing. In: Proc. Green Communic. Workshops at ICC. (2009)
16. Lange, C., et al.: Network element characteristics for traffic load adaptive network operation. In: Proc. ITG Symp. on Photonic Netw. (2012)
17. Feller, F.: Evaluation of centralized solution methods for the dynamic optical bypassing problem. In: Proc. ONDM. (2013)
18. Ungerman, J.: IP NGN backbone routers for the next decade (2010) http://www.cisco.com/web/SK/expo2011/pdfs/SP_Core_products_and_technologies_for_the_next_decade_Josef_Ungerman.pdf.
19. Alcatel-Lucent: New DNA for the evolution of service routing: The FP3 400G network processor (2011)
20. Wheeler, B.: EZchip breaks the NPU mold. Mircoprocessor Report (2012)
21. Cisco: CRS Carrier Routing System Multishelf System Description. (2011)
22. Cisco: CRS 16-slot line card chassis performance route processor data sheet (2012)
23. Cisco: CRS 1-port 100 gigabit ethernet coherent DWDM interface module data sheet (2012)
24. Cisco: CRS 2-port 40GE LAN/OTN interface module data sheet (2013)
25. Cisco: CRS 4-port 40GE LAN/OTN interface module data sheet (2012)
26. Cisco: 100GBASE CFP modules data sheet (2012)
27. Cisco: Pluggable optical modules: Transceivers for the cisco ONS family (2013)
28. Cisco: 40GBASE CFP modules data sheet (2012)
29. Cisco: ONS 15454 100 Gbps coherent DWDM trunk card data sheet (2013)
30. Cisco: ONS 15454 40 Gbps CP-DQPSK full C-band tuneable transponder card data sheet (2012)
31. Finisar: Stand alone 1RU EDFA (2012)
32. Cisco: Enhanced C-band 96-channel EDFA amplifiers for the cisco ONS 15454 MSTP data sheet (2012)
33. Cisco: ONS 15454 DWDM Reference Manual – Appendix A. (2012)
34. Cisco: ONS 15454 Hardware Installation Guide – Appendix A. (2013)
35. Rizzelli, G., et al.: Energy efficient traffic-aware design of on–off multi-layer translucent optical networks. Comput. Netw. **56**(10) (2012)
36. 80 PLUS: Certified power supplies and manufacturers – Cisco
37. Cisco: CRS Carrier Routing System 16-Slot Line Card Chassis System Description. (2012)
38. Malladi, K.T., et al.: Towards energy-proportional datacenter memory with mobile DRAM. In: Proc. ISCA. (2012)
39. Epps, G., et al.: System power challenges (August 2006) http://www.cisco.com/web/about/ac50/ac207/proceedings/POWER_GEPPS_rev3.ppt.
40. Vishwanath, A., et al.: Adapting router buffers for energy efficiency. In: Proc. CoNEXT. (2011)
41. EZchip: NP-4 product brief (2011)
42. Ungerman, J.: Anatomy of Internet routers (2013) http://www.cisco.com/web/CZ/ciscoconnect/2013/pdf/T-VT3_Anatomie_Routeru_Josef-Ungerman.pdf.
43. Mandviwalla, M., et al.: Energy-efficient scheme for multiprocessor-based router linecards. In: Proc.SAINT. (2006)
44. Le Rouzic, E., et al.: TREND towards more energy-efficient optical networks. In: Proc. ONDM. (2013)
45. Orlowski, S., et al.: SNDlib 1.0–Survivable Network Design Library. Netw. **55**(3) (2010) 276–286
46. Uhlig, S., et al.: Providing public intradomain traffic matrices to the research community. SIGCOMM Comput. Commun. Rev. **36** (2006)