A Quantitative Survey of the Power Saving Potential in IP-over-WDM Backbone Networks

Ward Van Heddeghem, Bart Lannoo, Didier Colle, *Member, IEEE*, Mario Pickavet, *Senior Member, IEEE*, and Piet Demeester, *Fellow, IEEE*

Abstract—The power consumption in Information and Communication Technologies (ICT) networks is growing year by year; this growth presents challenges from a technical, economic and environmental point of view. This has lead to a great number of research publications on 'green' telecommunication networks. In response, a number of survey works have appeared as well. However, with respect to backbone networks most survey works (a) do not allow for an easy cross-validation of the savings reported in the various works, (b) nor do they provide a clear overview of the individual and combined power saving potential. Therefore, in this work we survey the reported saving potential in IP-over-WDM backbone telecommunication networks across the existing body of research in that area. We do this by mapping more than 10 different approaches to a concise analytical model, which allows us to estimate the combined power reduction potential.

Our estimates indicate that the power reduction potential of the once-only approaches is $2.3 \times in$ a Moderate Effort scenario and $31 \times in$ a Best Effort scenario. Factoring in the historic and projected yearly efficiency improvements ("Moore's law") roughly doubles both values on a 10-year horizon. The large difference between the outcome of the Moderate Effort and Best Effort scenario is explained by the disparity and lack of clarity of the reported saving results, and by our (partly) subjective assessment of the feasibility of the proposed approaches. The Moderate Effort scenario will not be sufficient to counter the projected traffic growth, although the Best Effort scenario indicates that sufficient potential is likely available. The largest isolated power reduction potential is available in improving the power associated with cooling and power provisioning, and applying sleep modes to overdimensioned equipment.

I. INTRODUCTION

Power consumption in backbone telecommunication networks is still growing — The global amount of Internet Protocol (IP) traffic is growing every year. While this growth is gradually slowing down from an earlier Compound Annual Growth Rate (CAGR) of 100% (about 10 years ago) to an estimated CAGR of around 20–30% currently, this reduced growth still outperforms the annual 13% efficiency increase of new telecommunication equipment in the backbone network [1]. As can be seen in Fig. 1 this creates a so-called 'energy gap', and as such, the power consumed by telecom backbone network devices continues to increase year by year. This presents issues both from an economic (reducing the energy cost), technical (reducing the associated heat dissipation) and environmental (reducing the carbon footprint) point of view. And while today the power consumption in backbone networks



1

Fig. 1. Traffic growth and new equipment energy efficiency improvement in the IP backbone, normalized to the year 2010. As the traffic is growing faster than the efficiency improvement, this creates an 'energy gap'. Traffic estimations based on Cisco forecasts as detailed in Section IV-C, backbone efficiency ([C/P]=Gbps/W) improvement at +13% per year from Section IV-D.

makes up only about 8% of the total operator network power consumption (which includes the wired access, mobile access and backbone network) [2], with the expected increase of traffic volume, high growth rates in the backbone's energy consumption are expected, potentially even overtaking the access network's consumption¹[3]. The three issues mentioned above—economic, technical and environmental—are reflected in the increasing number of publications and research on this 'green' networking topic by academia, industry and governmental bodies alike [4].

Most 'classical' surveys of power saving approaches do not list and quantify the power saving potential — The research on 'green' Information and Communication Technologies (ICT) networks can be categorized in two broad categories: (a) estimating (current and future) network power consumption on the one hand; and (b) proposing and evaluating novel techniques to reduce the power consumption on the other hand. Especially for the latter, there are a number of surveys available that list and categorize power saving approaches. Notable ones focusing (partly) on optical backbone networks include those from Zhang et al. [5], Bolla et al. [6] and Bianzino et al. [7] as listed in Table I. We discuss these works in slightly more detail in Section II. Categorization of the power saving

The authors are with the Department of Information Technology (INTEC) of Ghent University - iMinds, Gaston Crommenlaan 8, 9050 Gent, Belgium

¹The reason is that the power consumption in wired access networks is proportional to the number of connected subscribers, while the consumption in the backbone network is proportional to the traffic volume [2].

approaches in the different survey works is done using a variety of different criteria, such as the area of application (e.g. circuitlevel versus network-level) or the considered time scale (e.g. sleep mode during operation versus energy-efficient network design in the planning phase). While these surveys certainly are worthwhile to make sense of the growing body of publications on this topic, it is a striking observation that none of the above three works list the power saving potential reported by (most of) the works they survey in a way that allows an easy cross-comparison. This is curious, as it would seem that the estimated (order of) energy saving potential is one of the main relevant outcomes in most of the surveyed works. Furthermore, such an inventory would give an idea of the relative power saving potential of the different approaches, and the consensus across the research community. This would allow-at least as a first step-to assess which techniques are most promising and should be prioritized from an energy saving point of view; a complete assessment would require consideration of other aspects such as implementation cost and associated operational issues.

Existing works that do survey the (combined) power saving potential have some shortcomings — To be fair, a number of works exist that do map the power saving potential of different power saving techniques, and/or their combined power saving potential. We have listed them in Table I. However, we think each of these works suffers from a number of different shortcomings. Specifically, we found that some studies either (a) lack sufficient references or traceability of the quoted savings ([8], [2]), (b) provide an assessment that is too rough ([9]), or (c) do not provide a combined saving potential ([11], [12]). Finally, while GreenTouch's Green Meter [10] does address the above issues, the limited number of referred works does not help the reader to get an idea on the agreement of the given power savings across the research community outside of GreenTouch. We discus the above cited works in more detail in Section II.

We address the above shortcomings based on an analytical model to survey and quantify the power saving potential in backbone networks — Our earlier work [13] was a first response to the various issues identified above; we complete that initial—but incomplete, as it did not contain a consistent and methodological survey, and no saving potential was calculated proposal in this study. Our methodology is based on a concise analytical model to estimate the power consumption of a backbone network². A comparable model was first used by Baliga et al. [3] and later formalized by Kilper et al. [1]. In [16], we have used a similar analytical formulation in our study on representative power consumption values for backbone equipment. It was not used to survey the power saving potential however. In essence, the total network power consumption is estimated by multiplying the power rating (W/Gbps) of all relevant backbone network equipment with the estimated total amount of traffic in the network. A number of correction factors are introduced in the above multiplication to account for the impact of traffic protection, overprovisioning, cooling power consumption, etc. In this study, we survey and map different energy saving approaches to each of the factors in the analytical model. This has the advantage that the interdependence of the various power saving approaches can be reduced or assessed more easily. Also, the model allows us to apply the impact of specific power saving approaches to specific equipment (such as only to IP routers, and not to transponders). Finally, as the model is based on a set of simple multiplications, we can rather easily calculate the combined impact of all power saving approaches.

Key contributions of this paper — Based on our model described above, we evaluate the power reduction potential for over 10 different approaches applied to an IP-over-WDM backbone network. The baseline of the equipment power efficiency is set at the year 2010. The key novelties of our work with respect to earlier publications are the following.

- First, we provide a formalized methodology to assess the combined power saving potential of various approaches. This assures a large amount of transparency to the power reduction calculations.
- Second, we categorize the power saving approaches by the multiplication factors in our model, and provide a tabulated survey of the most relevant publications for each approach, including the reported power saving potential. This provides the reader with an insight on the consensus of the values reported by different researchers.
- Third, for each approach, we infer a power reduction factor both for a Moderate Effort scenario and a Best Effort scenario. This provides the reader with a plausible range of the saving potential.

Limitations of this paper — We should not be mislead by the analytical nature of our methodology to overestimate the accuracy of the final results. As with all models, it has only limited capability in capturing and representing the actual situation it models. Therefore, it makes sense to specifically point out the scope and limitations of this work and approach. First, we focus on the electricity consumption of backbone networks only. We do not consider the embodied energy; the embodied energy is the sum of all energy required to produce equipment. Nor do we consider the opportunities (and issues) associated with using renewable energy sources to power network nodes. The wired and mobile access network are out of scope, as including these would make for an unreasonably large survey; in addition they would require separate models. Data centers are also considered out of scope, although we do briefly touch upon them when discussing caching (Section IV-C).

Second, our baseline is an IP-over-WDM network in the year 2010 (with a projection to the year 2020). As such, we do not take into account legacy network equipment and intermediate transport technologies (such as SONET/SDH, or OTN). This is an important limitation; we provide some discussion on the applicability of our findings and potential impact of legacy

²An alternative technique to estimate the network power consumption is to use a top-down approach, where the total telecom operator power consumption is trimmed from non-relevant contributors (such as data centers, telecom offices) to give only the network relevant contribution. This has been used by e.g. Malmodin et al. [14] and Lambert et al. [15] for estimating the worldwide telecom operator network electricity consumption.

	TABLE I		
AN OVERVIEW OF SURVEYS ON ENERGY EFF	ICIENCY IN (BACKBONE) TELECOMMUNICATION NETWO	ORKS

Publication	Scope	Reported power savings
Zhang, 2010 [5]	Optical networks (fixed access, metro, core)	-
Bolla, 2011 [6]	Wired networks	-
Bianzino, 2012 [7]	Wired networks	-
Roy, 2008 [8]	Wireline and wireless	Energy saving potential: wireline $\approx 40\%$, wireless $\approx 60\%$
Lange, 2011 [2]	Home networks up to backbone networks	Energy saving potential in 2007: backbone 65%, total network 56%
Parker, 2011 [9]	Photonic telecom networks	Energy efficiency improvement: $1000 \times$
Kilper, 2012 [4]	Optical transmission networks	-
Green Meter, 2013 [10]	Mobile access, wireline access, core network	Energy efficiency improvement (2010–2020): mobile access $1043\times$, wireline access $449\times$, core $64\times$
Le Rouzic, 2013 [11]	Optical networks (fixed access, metro, core, data centers)	Individual savings only, no consolidated savings
Le Rouzic, 2013 [12]	Metro, core, data centers	Individual savings only, no consolidated savings

network equipment in Section V-E.

Third, our breakdown of network power consumption in a number of factors restricts the scope of the power savings approaches we can reasonably capture. For example, we do not specifically identify approaches such as 'multi line rate' versus 'single line rate' transmission, or Elastic Optical Networks (EONs). However, we do think that we capture the majority of approaches reported in the literature.

Fourth, this work remains a survey, and is not an evaluation over a standardized benchmark or baseline. The material we survey is from publications with a wide range of quality control; nonetheless we have tried to use journal papers instead of conference papers when available, and tried to give more weight to results from publications that appeared to be more thorough. We perform a best-effort estimation of the saving potential of various approaches. Evaluating most of the presented techniques on a consistent baseline scenario (topology, traffic matrix, equipment power consumption values) would be the next logic step to do, but is out of scope of this paper.

Organization of this paper — For the remainder of this paper, we first describe the earlier mentioned related work in more detail in Section II. We outline our methodology and the associated analytical power model in Section III, before using this model in Section IV to survey a number of power reduction approaches across our baseline. Finally, in Section V we present the combined power saving potential, discuss its sensitivity to a number of model parameters such as the network hop count, and compare them to the Green Meter results [10]. In Section VI we summarize our findings and complement it with a number of recommendations for the research community, network operators, vendors and policy makers.

II. RELATED WORK

This section discusses the main works that have been cited in the preceding section. The first paragraph covers general survey papers. The second paragraph covers works that do quantify the isolated and combined power saving potential in backbone networks, and highlights some of their shortcomings. The main works are listed in Table I.

In the work by Restrepo et al. [17], power reduction approaches are broken down into three levels depending on the area of application: (i) on circuit level (such as the use of dynamic voltage or frequency scaling techniques), (ii) on equipment level (e.g., replacing components by their counterpart in the optical domain) and (iii) on network level (network planning for efficiency, e.g. by using optical bypass). In [5], Zhang et al. consider the following techniques to improve the energy efficiency in the backbone network: (i) selectively turning off network elements, (ii) energy-efficient network design, (iii) energy-efficient IP packet forwarding, and (iv) green routing. The power savings for the surveyed works are listed as either 'low', 'medium' or 'high'. In [18], and the earlier survey paper [6], Bolla et al. classify the approaches as either: (i) re-engineering (more energy-efficient network elements, e.g. replacing electronics by optics where possible), (ii) dynamic adaptations (scale power consumption with actual load, e.g. dynamic voltage or frequency scaling), and (iii) sleeping/standby (drive unused network devices to low standby modes). The power savings for the surveyed works are listed throughout the text, but not cross-compared. Bianzino et al. identified in [7] four branches of green networking research, being (i) adaptive link rate (including sleep mode and rate switching), (ii) interface proxying, (iii), energy-aware infrastructure, and (iv) energy-aware applications. The power savings for the surveyed works are not listed. As is clear from this brief survey, many different approaches have been used to categorize power saving approaches. All these categorizations come with their own merits and drawbacks. We feel that categorizing the different approaches according to the factors in a concise analytical model allows for a more insightful estimation to quantitatively assess the potential power savings.

The first work, to our knowledge, that does provide a quantified breakdown of various approaches to reduce the power in telecommunication networks is by Roy [8]; however, the work is rather economical on references. The study by Lange et al. [2] is very holistic and complete in that it does a thorough assessments of the power consumption across the various network segments, expected trends, and opportunities to increase the energy efficiency. However, the quantified saving potential in the backbone is only split up in two high level approaches (technology progress/energy aware systems,

and load adaptive network operations), and the associated savings (55% and 20%, respectively) are hard to verify or trace back. The two papers by Kilper [1], [4] span a similar range of holistic topics as the work by Lange, however the saving potential of various approaches is not explicitly quantified. The work by Parker and Walker [9] was the initial inspiration for our study. Next to proposing an absolute energy efficiency metric $(dB\varepsilon)$ for any ICT system, it presents a notable first effort to provide a synoptic analysis of the performance of 10 different techniques to achieve a 1000-fold reduction in the power consumption of future photonic networks; however, their estimates might be rather optimistic. The Green Meter [10] by the GreenTouch consortium is probably closest in spirit to our work. For backbone networks it identifies and quantifies the power reduction potential of seven approaches, and the combined overall reduction factor. While the fact that some of the approaches are evaluated on a common benchmark presents a rare (but deeply needed) advantage over the other works, this is also stems from the fact that most of the cited works refer to only a few partners in the GreenTouch consortium; as such, the consensus over its findings across the research community outside of GreenTouch cannot be assessed by the reader. The two works by Le Rouzic et al. [11], [12] present an interesting overview of the different backbone saving potentials as identified in the European TREND network of excellence. However, the quoted savings are mainly isolated values; their interplay and the combined saving potential is not given.

Finally, it is worth mentioning the work by Masanet et al. [19]. They estimate the energy use and efficiency potential of U.S. data centers using a methodology very similar to our study. They use a bottom up analytical modelling approach to capture and categorize the total data center electricity demand, and discuss and evaluate a number of efficiency measures. These measures are then presented as reduction factors which can be multiplied to compute the combined electricity saving potential.

III. METHODOLOGY

Our methodology to assess the impact of efficiency improvements in backbone networks is based on a concise analytical model, which is described by the general form in Eq. (1). It expresses the network-wide power consumption associated with an equipment type (such as an IP router) based on a number of factors. Basically, the power rating of the considered equipment (W/Gbps) is multiplied with the amount of traffic in the network (Gbps) and the number of network hops. Three correction factors account for the power consumption associated with general overhead, traffic protection, and overprovisioning. A hop count correction factor is relevant to correctly account for the hops in either the IP layer, or the Wavelength Division Multiplexing (WDM) layer. By considering reduction factors for each of the parameters in Eq. (1), we can derive the total reduction potential across an energy-optimized network.

$$P_X = \eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot (H + c_x) \cdot W_x \tag{1}$$

where

- P_X total network power consumption for equipment type x (W);
- η_{eo} external overhead factor, e.g. cooling;
- η_{pr} protection factor;
- η_{op} overprovisioning factor;
- T total traffic in the network (Gbps);
- *H* average hop count in the respective network layer;
- c_x hop count correction factor (0 or 1)
- W_x weighting factor of equipment x, based on its power rating $\frac{P_x}{C_x}$ (W/Gbps) multiplied with an equipment dependent correction factor.

This short overview of our methodology provides us with the required background to outline the (more detailed) remainder of this section.

We first define our baseline, i.e. the reference network scenario over which we will evaluate the different power saving approaches; the equipment power rating values will serve as baseline power weighting factors for the different equipment types (Section III-A). We then discuss the analytical model in more detail and expand on the general form given above (Section III-B). With this established, we detail how this model is used to calculate the impact of efficiency improvements from different contributors on the complete backbone power consumption (Section III-C). Following a short discussion on different ways to express power savings (Section III-D), we introduce two savings scenarios which will provide a lower bound and upper bound for the likely saving potential in backbone networks (Section III-E).

A. Baseline scenario and reference power rating values

Our baseline architecture is an IP-over-WDM network. As shown in Fig. 2, we consider IP routers (which include line cards), long-haul transponders (labelled TXP), Optical Cross-Connects (OXCs), and Optical Line Amplifiers (OLAs). The



Fig. 2. Conceptual model of our generalized IP-over-WDM network architecture. It shows a bidirectional working path (solid lines) and backup path (dashed lines) under a 1+1 protection scheme. (LC = Line Card, TXP = Transponder, OXC = Optical Cross-Connect, OLA = Optical Line Amplifier)

IP routers switch the IP traffic in the IP layer. The longhaul transponders transmit and receive the optical signal over dedicated wavelengths in fibers. The transponder capacity is 10G, and the fiber multiplexes 40 wavelengths. OXCs provide optical switching capabilities in the WDM layer, by adding and dropping the wavelengths in the different network nodes as required. OLAs are required typically every 80 km and amplify all wavelengths in a fiber. As there is a move from multiple stacked technologies to IP-over-WDM, we leave out other potential intermediate switching technologies, such as Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH), Optical Transport Networking (OTN) and Ethernet. Furthermore, we do not include the power consumption of 3R regenerators (required for optical channels spanning large distances, a typical value is over 1500 km), as their contribution to the total power consumption is marginal [16].

The power rating values we will consider for the four types of equipment are given in Table II. The power rating expresses the power per capacity, in W/Gbps, and is as such an (inverse) measure for the power efficiency of the considered equipment. We will use them as a basis for the weighting factors (also listed in Table II) in the analytical model, as we will explain in the next subsection. The power ratings are distilled from [16] and [20], and are homogenized across three properties. First, they are derived from typical power consumption values (as opposed to vendor rated power, which is typically higher and used for provisioning the power distribution infrastructure). Second, the power rating includes the power associated with the chassis and control overhead (as opposed to only the power for the functional component, which might lead to very low values for e.g. OLAs). Third, the power rating value is that for realistic filling levels of shelves and racks (as opposed to optimal power rating values when assuming maximally filled shelves/racks) [20].

Finally, these values are relevant for the year 2010, which is the reference year in our baseline scenario. The reason we choose the year 2010, is because that is the most recent year for which we have reliable power consumption values available for *all* backbone equipment considered, as provided by our earlier publication [16] which was dedicated to this.

 TABLE II

 Power rating values and weighting factors, reference year

 2010 (based on [16], [20])

Type	Power	Power rating	Weight
1)pe	10000	D/C	II/
		P_x/C_x	VV x
Core IP router (inc. line cards)	-	10.00 W/Gbps	10.00
			0.00
OXC (a)	-	0.46 W/Gbps	0.92
Transponder 10G	50 W	5.00 W/Gbps	10.00
OLA 80 km ^(b)	165 W	0.41 W/Gbps	2.06

^(a) Power rating 0.46 W/Gbps of an OXC node degree 3, for $40 \times 10G$ channels. ^(b) Power rating 0.41 W/Gbps of a bidirectional OLA for $40 \times 10G$ channels.

B. Analytical power model

The analytical power model that we will use is inspired both by the work from Baliga et al. [3], Kilper et al. [1] and the findings from our earlier work [16].

The total power $P_{BACKBONE}$ in an IP-over-WDM network can be given as the sum of the power consumption in the constituting layers:

$$P_{BACKBONE} = P_{IP} + P_{WDM} \tag{2}$$

$$= P_{IP} + \left(P_{OXC} + P_{TXP} + P_{OLA}\right), \quad (3)$$

with P_{OXC} , P_{TXP} , and P_{OLA} respectively being the total network power consumption of the OXCs, the WDM transponders, and the OLAs.

The power consumption of each equipment type is further given as:

$$P_{IP} = \eta_{eo} \cdot \frac{\eta_{pr}}{2} \cdot \eta_{op} \cdot T \cdot (H+1) \cdot \left(\frac{P_{ip}}{C_{ip}} \cdot 2\right) \tag{4}$$

$$P_{OXC} = \eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot H \cdot \left(\frac{P_{oxc}}{C_{oxc}} \cdot 2\right)$$
(5)

$$P_{TXP} = \eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot H \cdot \left(\frac{P_{txp}}{C_{txp}} \cdot 2\right) \tag{6}$$

$$P_{OLA} = \frac{\eta_{eo}}{2} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot H \cdot \left(\frac{P_{ola}}{C_{ola}} \cdot \left\lfloor\frac{\text{link length}}{80 \,\text{km}}\right\rfloor\right)$$
(7)

Note that we use upper case subscripts for the total network power consumption (e.g., P_{IP}), and lower case subscripts to denote a single equipment power consumption value (e.g., P_{ip}).

The external overhead factor η_{eo} accounts for the power consumption due to external cooling and facility overheads in telecom centers, with typically $\eta_{eo} \approx 2$. This value is not applicable to OLAs as they are typically deployed in dedicated outside cabinets without active cooling, so we apply a correction factor $\frac{1}{2}$ to η_{eo} in Eq. (7). The protection factor η_{pr} accounts for traffic protection, with $\eta_{pr} \approx 2$ for backbone networks using a 1+1 protection scheme (i.e., all traffic is routed twice on link-disjoint paths). As this is typically done in the WDM layer but not the IP layer (see further), again we apply a correction factor $\frac{1}{2}$ to η_{pr} in Eq. (4). The overprovisioning factor η_{op} accounts for the overprovisioning of the network capacity to deal with unexpected traffic spikes and future traffic growth. The traffic factor T gives the total amount of traffic in the network (in Gbps). The hop count His the average number of hops between processing elements in the respective layer. The reason we have (H+1) for the IP layer is because we also need to account for the client side capacity of the IP router, i.e., towards the access network³; for more details see [16]. The power rating factor $\frac{P_x}{C}$ expresses the average power per capacity (in W/Gbps) for a given equipment x, as listed in Table II. The factor 2 at the end of the equation accounts for the fact that for each hop the relevant node

³In [16] we actually used $(H + 1/\eta_{pr})$ instead of (H + 1). However, the slight increase in accuracy does not affect our results in any meaningful way, but does provide additional complexity to calculate the power saving potential. Therefore, we just use the more simple term (H + 1).

6

capacity is required at both the sending and receiving side⁴. The OLA power consumption is a function of the average link length, as an OLA is required every 80 km.

To further generalize on the four equations above, we will combine the factors between brackets and the two $\frac{1}{2}$ correction factors for IP routers and OLAs into a weighting factor W_x , resulting in the following general equation:

$$P_X = \eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot (H + c_x) \cdot W_x \tag{8}$$

with $c_x = 1$ for the IP layer, and $c_x = 0$ for the equipment in the WDM layer. Note that is exactly the general form of our analytical model as in Eq. (1). The weighting factors W_x are listed in Table II. For the OXCs, and transponders the weighting factor is simply twice the power rating value $\frac{P_x}{C_x}$. For the IP routers the weighting factor is equal to the power rating value. For the OLA contribution in Eq. (7), we calculated the weighting factor assuming an average link length of 800 km, which is a reasonable value for backbone networks. In Section V-B we will look at the impact of the link length on our efficiency improvement estimates, and show that it is almost negligible.

C. Calculating power savings

1) Savings for one equipment type: To estimate the savings for an equipment type x in a backbone network, we can model an improvement in each of the factors in Eq. (8) by inserting beta reduction factors (with $\beta \ge 1$).

$$P_{X,impr} = \frac{\eta_{eo}}{\beta_{eo}} \cdot \frac{\eta_{pr}}{\beta_{pr}} \cdot \frac{\eta_{op}}{\beta_{op}} \cdot \frac{T}{\beta_t} \cdot \left(\frac{H}{\beta_h} + c_x\right) \cdot \frac{W_x}{\beta_{pc}} \tag{9}$$

Each improvement factor β can be seen as an *approach* that acts independently to reduce the power consumption. The goal of the upcoming Section IV will be to determine feasible β values for a wide number of approaches.

Thus, the total power reduction $\beta_{x,tot}$ for the equipment type x of an improved network consuming $P_{X,impr}$ power compared to the referenced network consuming $P_{X,ref}$ is given by:

$$\beta_{x,tot} = \frac{P_{X,ref}}{P_{X,impr}} \tag{10}$$

$$= \frac{\eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot (H + c_x) \cdot W_x}{\frac{\eta_{eo}}{\beta_{eo}} \cdot \frac{\eta_{pr}}{\beta_{pr}} \cdot \frac{\eta_{op}}{\beta_{op}} \cdot \frac{T}{\beta_t} \cdot \left(\frac{H}{\beta_h} + c_x\right) \cdot \frac{W_x}{\beta_{pc}}}$$
(11)

Note that we can cancel out the factors η_{eo} , η_{pr} , η_{op} and T.

2) Savings across multiple equipment types: The total power consumption in an IP-over-WDM network is the sum of the power in the IP layer, the OXCs, the transponders and the OLAs. To calculate the total power reduction β_{total} across the

complete network, we now have:

$$\beta_{total} = \frac{P_{ref}}{P_{impr}}$$

$$= \frac{\sum_{x}^{equip} \left[\eta_{eo} \cdot \eta_{pr} \cdot \eta_{op} \cdot T \cdot (H + c_x) \cdot W_x \right]}{\sum_{x}^{equip} \left[\frac{\eta_{eo}}{\beta_{eo,x}} \cdot \frac{\eta_{pr}}{\beta_{pr,x}} \cdot \frac{\eta_{op}}{\beta_{op,x}} \cdot \frac{T}{\beta_{t,x}} \cdot \left(\frac{H}{\beta_{h,x}} + c_x \right) \cdot \frac{W_x}{\beta_{pc,x}}} \right]$$

$$(12)$$

Note that the β factors can be different for each equipment type, e.g. β_{pr} can be different for the IP routers and the transponders; therefore we have added an x index to each β factor.

The occurrence of the c_x term is unfortunate, as it makes the result dependent on the hop count H (i.e., we can not factor out all β factors). While for a given topology the hop count will depend on several aspects, such as the routing algorithm and link weights, a good ballpark number of H in a backbone network is 3–4 hops [1], [16], [21]. Therefore we will assume H = 3 when calculating the saving potential in the next section. In Section V-C we will look at the impact of changing the hop count H on our power improvement estimates, and show that it is rather limited.

D. Expressing power savings

In the above sections, we have modelled power improvements using a *reduction factor* $\beta \geq 1$. In contrast, by far the most common approach used in publications on energy efficiency in networks is to state the *savings percentage* γ (with $\gamma = 1 - \frac{1}{\beta}$). For example, a power reduction with a factor $\beta = 4$ corresponds to a savings percentage $\gamma = 75\%$. It might be interesting to note that another variation is to express the reduction factor in decibel⁵ (dB), i.e. on a logarithmic scale as $10 \log_{10} \beta$.

While both the reduction factor β and the savings percentage γ are mathematically interchangeable, the power reductions intuitively communicated by them is different. For example, when comparing the power savings $\gamma_1 = 80\%$ and $\gamma_2 = 90\%$, the latter appears to be only a slight improvement over the first one. While this is indeed true compared to the baseline (i.e., the original power consumption), it does somewhat conceal the fact that the 90% savings scenario consumes only half of the power of the 80% savings scenario. Given the historic and projected exponential growth of traffic in the backbone network (see Fig. 1), it is important to realize we need and are interested in significant power reductions. A power saving approach (or more likely, any combination of approaches) that can provide 90% savings instead of 80% savings allows for an extra doubling of the traffic while still consuming the same amount of power. In contrast, if we would have identified the power savings using a reduction factor $\beta_1 = 5 \times$ and $\beta_2 =$ $10 \times$ respectively, the factor 2 in difference would have been instantly clear.

⁴This approach implies that all traffic is bidirectional, i.e., that there is as much traffic from node A to node B, as from B to A. While this is a simplification, the approximation is adequate enough for our purpose.

⁵Expressing the power reduction in decibel has the advantage that improvements can easily and intuitively be summed (instead of multiplied), and is for example for this reason used in [9].

Therefore, in the upcoming Section IV we will express our estimates of the power reduction potential of the various approaches using β reduction factors (≥ 1). However, when we quote and list existing works (such as in Table VI) we will cite the savings as a percentage γ so that the values given can easily be verified in the referenced works. For the reader's convenience, we have listed a set of savings percentages and the corresponding reduction factors in Table III; if required, this can be consulted when going through the subsequent sections.

 TABLE III

 Conversion table for the reader's convenience

Savings percentage γ	Reduction factor β
0%	$1 \times$
10%	1.11×
20%	$1.25 \times$
30%	$1.43 \times$
40%	$1.67 \times$
50%	$2.00 \times$
60%	$2.50 \times$
70%	3.33×
80%	$5.00 \times$
90%	10.00×

E. The Moderate Effort and Best Effort scenario

In our analysis in the upcoming Section IV, we will distinguish between two scenarios to model the power savings that are possible in backbone networks.

- In the Moderate Effort scenario we use small power reduction estimates, corresponding to solutions that are relatively feasible from either a technical or operational point of view. Where the power reduction associated with an approach is unsure, we will consider a likely lower bound that we will model in this scenario.
- In the **Best Effort scenario** we use larger, more aggressive power reduction estimates, corresponding to solutions that would be more challenging to implement from a technical, operational or cost perspective. Where the power reduction associated with an approach is unsure, we will consider a likely upper bound that we will model in this scenario.

Combined, both scenarios provide a range for the achievable power savings potential in backbone networks.

IV. APPROACHES TO SAVE POWER

In this section, we discuss several approaches to reduce the power consumption in backbone networks. We do so using the parameters in Eq. (1). Fig. 3 gives an overview of the different power saving approaches that we will discuss, and how they map to the various factors.



Fig. 3. Overview of power reduction approaches mapped to the general form of our analytical model

A. External overhead factor η_{eo}

The external overhead factor η_{eo} accounts for power consumption associated with external cooling and facility overheads in telecom centers. This overhead is commonly characterized by, and also commonly known as, the Power Usage Effectiveness (PUE). The PUE is the ratio of the total amount of power consumed to the useful power consumed⁶, and typically has a value of about 2 [24]. In this specific case, this means that for each Watt consumed by useful equipment, such as servers and switches, an additional Watt is consumed through external overhead. In highly optimized and efficiently cooled data centers (much) lower PUE values are possible⁷, but this is not yet commonplace. On average, this overhead is made up of two main contributing components [27][28]: cooling and air conditioning, and efficiency losses in power provisioning (see Fig. 4); the contribution of switchgear and lighting is only minor.

We discuss three approaches to reduce the external overhead factor η_{eo} : more efficient cooling systems, high-temperature chips, and more efficient power provisioning.

More efficient cooling — A first approach is to increase the efficiency of the premises cooling, i.e. to provide the same degree of premises cooling effect while using less electrical energy. This is a particularly hot topic in data center research, and a wealth of publications on this topic is available. Good

⁶While the PUE concept is elegant and simple, in practice there are many intricacies that make it easy to result in different values. To illustrate, while the 2008 white paper that documents the PUE methodology contained only 9 pages [22], the version from 2012 consists of 83 pages [23].

⁷For example, Google states to have reached an annualised average PUE across all their tracked data centers of 1.14 by the end of 2011 (with a minimum value of 1.11 and a maximum value of 1.21) [25]. This value is even more impressive, as they claim that their PUE calculation is more comprehensive than what is done by other players, by accounting also for overhead sources that are typically omitted, such as data center offices and site substation losses. According to the 2012 Data Center Survey by the Uptime Institute[26], the PUE reported by its participants averages between 1.8 and 1.89.



Fig. 4. Rough distribution of the external overhead power among its main contributors (source: [27], [28], [8])

overviews are available in [29], [30] and [31]. Examples of such approaches include hot aisle/cold aisle (to avoid mixing both cold and hot air), free cooling (using cold outside air, if ambient temperature and humidity permit), and rack liquid cooling (to improve heat transfer).

High-temperature chips — A straightforward alternative to reduce cooling power is to cool less [30]. For example, it has been observed by a notable study from Google [32] that contrary to popular belief, hard disks do not become less reliable when running at higher temperatures. However, this is only possible up to a certain limit. An approach taking the concept beyond this limit is that of research into the high-temperature operation of integrated circuits [9]. Such an approach further reduces the need for cooling, and as such brings down the external overhead factor η_{eo} .

More efficient power provisioning — Power provisioning accounts for roughly 20% to 30% of the external overhead power [28], [8], mainly through inefficiencies in Uninterruptible Power Supply (UPS) units and Power Distribution Units (PDUs) [27]. While the efficiency of a UPS unit can be around 90% at maximum load, its efficiency drops off steeply when lightly loaded—which is very often the case [33]. One reason for lightly loaded provisioning equipment is that such equipment is deployed based on nameplate power ratings (i.e. vendor indication of the maximum power drawn, which is used for dimensioning the power supply systems) of the ICT equipment, which can be substantially higher than the actual peak power [34]. Right-sized provisioning of equipment to the actual peak power of the ICT equipment would reduce the overhead power consumption.

To asses the total overhead power saving potential, we could try to assess the individual potential of each of the above three methods. With cooling comprising typically more than 60% of this overhead ([28], [8]), the biggest gains can probably be achieved there. The study by Roy [8] gives an indication⁸. It estimates that implementing cooling best practices and supplemental high density cooling can save 44% on the required cooling power. Optimizing the power provisioning by replacing legacy rectifiers with new generation rectifiers (to increase the peak efficiency), using DC-powered

IT equipment, and employing DC ECO mode (to improve rectifier efficiency at lower loads) can save 71% on the power provisioning losses. Combined, this reduces Roy's central office PUE value from 2.14 to 1.63, wich is a reduction of about $1.3 \times$ (or 23% savings). However, this value depends on the original PUE value. Therefore, we will go for a more pragmatic approach and lump all approaches together in the PUE factor, and explore how this factor can be improved. To do so, we require (a) a good estimate on the baseline PUE value for telecom network infrastructure in 2010, and (b) a realistic estimate on the improved PUE achievable at these premises.

While there have been a number of studies on the PUE values of data centers (such as [24], [26]), little information is available on a baseline PUE for telecom network infrastructure. A first approach would be to just use the average data center PUE value. However, we see a few reasons why the average PUE for data centers might not be representative for those of telecom network infrastructure. First, the PUE of telecom network infrastructure might be worse than those for data centers as network equipment typically has a longer operating lifetime and requires high uptimes, so an overhaul of supporting networking infrastructure to improve the cooling overhead is less easy to occur. On the other hand, the PUE of telecom network infrastructure might be better because it does not suffer from the effect of 'in-house' data centers. In-house data centers are facilities owned and operated by companies whose primary business is not computing, and as such efficient cooling is not much of a concern, resulting in relatively bad PUEs. These in-house data centers dominate the data centers in total electricity use and this worsens the average PUE of data centers [35]. In Table IV we list some indicative, public values to determine a reasonable value for telecom network infrastructure PUE in 2010. We have also included some data center values from telecom operators; these values are typically used by operators to highlight their energysavings efforts, and it seems reasonable to use them as a lower bound for their network infrastructure PUEs. In addition to the limited set of public values in the table, we should note that in private communication with operators, one operator reported an average around 1.5–1.6, while according to another operator values up to 2.0 are not an unreasonable average. Taking all of the above into account, an average baseline PUE value of 1.7 for telecom network infrastructure in 2010 seems a reasonable assumption⁹. Note that this is not too different from an average data center PUE value of 1.88 for the year 2010 [36].

For determining a feasible *improved PUE* value, even less public data is available. The work by Le Masson [41] applies partly to telecom equipment, and describes an experimental setup of a wall structure to improve the free cooling efficiency. The work reports a 6-month average PUE value of 1.28; however, the overhead power does not include any power provisioning, so the value is rather optimistic. The theoret-

⁹For our analytical model in Section III-B, we approximated the external overhead factor with $\eta_{eo} = 2$ (instead of 1.7), which we used to apply a $\frac{1}{2}$ correction for OLA equipment in Eq. (7). This difference does not affect the end results in Section V in a meaningful way.

TABLE IV INDICATIVE NATIONAL OPERATOR PUE VALUES (INCLUDING FOR THEIR DATA CENTERS (DCS)) FOR DETERMINING TELECOM NETWORK INFRASTRUCTURE (TNI) PUE VALUES

Operator/source	Year	DC PUE	TNI PUE
A new Orange DC [37]	2011	1.3	-
Average Telefonica DC [38]	2013	2.4	-
A new Telefonica DC [38]	2013	1.3	-
Wireline central office [8]	2008	-	2.14
Deutsche Telekom average [39]	2005	-	1.75
Deutsche Telekom average [39]	2009	-	1.53
Deutsche Telekom average [40]	2012	-	1.48
Expert interview national operator X	2013	-	1.5-1.6
Expert interview national operator Y	2013	-	1.8-2.0

ical lower bound of a PUE is 1.0, and highly optimized data centers do reach values as good as 1.12 in 2012 [25]. However, data centers can be built at locations suitable for very efficient cooling (such as close to a river, or at cold geographical locations), whereas telecom central offices are more constrained by the network topology and have been historically put e.g. in the middle of a city. In addition, the issue of legacy network equipment and the required high levels of availability, as mentioned above, makes upgrades to the supporting network infrastructure more complicated than is the case for data centers. So, it seems unlikely that in the medium term the PUE value of telecom network infrastructure can be better (i.e., lower) than 1.2.

Using both the above derived baseline PUE of 1.7 and the improved PUE of 1.2, we get an improvement factor of $1.4 \times$. We assume a slightly lower Moderate Effort reduction factor of $1.3 \times$ (or 23% savings); note that this is in line with the estimation from Roy [8] discussed higher up. For the Best Effort reduction factor we assume $2 \times$ (or 50% savings), as the baseline PUE might be higher and/or the improved PUE might be lower than what we derived (with the former situation seeming more likely). Since the PUE is not applicable for OLAs (they are deployed in dedicated outside cabinets without active cooling), the final **Moderate Effort reduction** factor becomes $1.27 \times$ (or 21% savings) and the **Best Effort reduction factor** becomes $1.85 \times$ (or 46% savings), as shown in Table V.

B. Protection factor η_{pr}

The protection factor η_{pr} accounts for the additional power consumption due to traffic protection. Traffic protection is typically employed in backbone networks to achieve high reliability to meet costly Service Level Agreements (SLAs). A typical approach in backbone networks is to provide Shared Path Protection (SPP) protection for the IP layer, and to employ a 1+1 protection scheme for lower layers (also referred to as Dedicated Path Protection (DPP)) [48], whereby for each demand between a source and destination node two link-disjoint WDM connections are set up. Some operators adopt an alternative strategy where, in essence, they deploy the backbone network twice with a certain equipment mix

TABLE V REDUCTION FACTOR FOR EXTERNAL OVERHEAD

	Mor	eft. co	Jing (9)	hip boner	bio4.
Moderate Effort reduction					
IP	-	-	-	$1.30 \times$	
Opt. Swit	-	-	-	$1.30 \times$	
Transponder	-	-	-	$1.30 \times$	
OLA	-	-	-	$1.00 \times$	
Weighted total ^(a)				1.27×	(=-21%)
Best Effort reduction					
IP	-	-	-	$2.00 \times$	
Opt. Swit	-	-	-	$2.00 \times$	
Transponder	-	-	-	$2.00 \times$	
OLA	-	-	-	$1.00 \times$	
Weighted total ^(a)				1.85×	(=-46%)

^(a) See Section III-C for calculation details.

^(b) These columns are empty as we estimated the *overall* PUE reduction instead. See text for more details.

using systems from different vendors, with the goal of having diversification over different vendors. Each of the two networks is dimensioned for full capacity, but consequently running at 50% utilization, in order to be able to overtake the whole traffic volume. Both of the above approaches result in a baseline protection factor of roughly $\eta_{pr} = 2$, meaning that due to protection the power consumption is doubled compared to a non-protected network.

We consider the following approaches to bring down the protection factor: avoiding dualling, sleep modes, and differentiated QoP. An overview of the main works cited in the subsequent paragraphs is given in Table VI.

Avoiding dualling — Parker et al. argue in [9] that energy savings of up to 50% of the total network power consumption could be achieved by avoiding the above described network dualling. This would be made possible through increased reliability of network devices, systems and subsytems, and increased software-defined operation of many significant network functionalities. However, this might be an overly optimistic estimation: it is, for example, unclear how more reliable equipment would address the issue of cable cuts. A more viable approach could consist of using more passive forms of protection (passive in the sense of not consuming power). While this would be hard for electrically switched networks (i.e., IP, SDH/SONET and OTN), this is feasible in the optical layer through employing 1x2 passive couplers that duplicate the lightpath from a single transponder (although in that case the service is not protected against transponder failures).

Sleep modes and PAR — An alternative approach to reduce the protection factor η_{pr} is to put protection equipment that is serving backup links into a low-power sleep mode. For realistic application, this would require fast sleep and wakeup times. Putting protection equipment to sleep is relatively straightfor-

TABLE VI						
PROTECTION FACTOR RELATED ENERGY SAVINGS REPORTED IN PUBLICATIONS						

App.	Source	Savings	Remarks	Justification			
Avoidi	ng dualling Parker, 2011 [9]	50%	Probably overly optimistic for all network layers. Viable approach for WDM layer.	-			
Sleep	modes and Power Aware Ro	outing (PAR)					
•	Cavdar, 2010 [42]	30-40%	WDM layer only. ILP optimization on COST239 network. SPP. For breakdown over impact of sleep mode and PAR, see Fig. 5.	Section V.C and Fig. 4 in [42]. Compar- ing <i>EASP</i> with <i>SBP</i> at 250 Gbps (low load) and 750 Gbps (high load).			
	Jirattigalachote, 2011 [43]	30-40%	WDM layer only. Simulation of dynamic content provisioning on COST239 and USNET. DPP. For breakdown over impact of sleep mode and PAR, see Fig. 5.	Fig.3 (a) and (b) in [43]. Comparing <i>EA-DPP Sleep</i> with <i>SP-DPP Total</i> .			
	Coiro, 2011 [44]	35%	WDM layer only. Exploits adaptation to daily traffic variations, by turning off links. ILP optimization and heuristic on 18-node random generated network. SPP. Part of the savings are attributed to 'overprovisioning typical of transport networks'.	Fig. 5 in [44].			
	Bao, 2012 [45]	45-65%	WDM layer only. Heuristic on COST239 and USNET. SPP. For breakdown over impact of sleep mode and PAR, see Fig. 5.	Fig 5 and Fig 6 in [45]. Comparing <i>EASPP</i> curve (referred to as <i>PASPP</i> in text) with 100% baseline.			
	Musumeci, 2013 [46]	36-45% DPP 29-37% SPP	IP-over-WDM. ILP optimization. Savings reported for NSFNET, but claiming to be similar for COST239. Contribution of PAR to savings unclear. Both DPP and SPP. Upper range of savings for higher traffic volumes.	Table 3 in [46]. Comparing for scaling factor $f = 1$ and $f = 10$ the (a) <i>DPP</i> sleep-mode result with <i>DPP all-ON</i> , and <i>SPP sleep-mode</i> result with <i>SPP all-ON</i> .			
Differe	Differentiated Quality of Protection (OoP)						
	Lopez, 2013 [47]	10–20%	WDM layer only. Simulation on Telefónica Spanish core network, for three scenario's with a different QoP class mix. Upper range of savings for higher traffic volumes.	Table 4 in [47], savings (roughly averaged) for traffic 1.56 Tb/s and 23.43 Tb/s.			

Acronyms: EASP: Energy-Aware Shared Backup Protection, CPE: Customer Premises Equipment, DPP: Dedicated Path Protection, EA-DPP: Energy Aware Dedicated Path Protection, EASP: Energy-Aware Shared Backup Protection, EASPP: Energy-Aware Shared Path Protection, GPON: Gigabit Passive Optical Network, ILP: Integer Linear Programming, PAR: Power Aware Routing, PASPP: Power-Aware Shared Path Protection, SBP: Shared Backup Protection, SP-DPP: Shortest Path Dedicated Path Protection, SPP: Shared Path Protection.



Fig. 5. Normalized (to baseline) power consumption of several power saving algorithms as reported for the COST239 network by (a) Cavdar [42], (b) Jirattigalachote [43], and (c) Bao [45]. The algorithms have been tagged as either sleep (where port-based equipment is put to sleep), Power Aware Routing (PAR) (where OLAs are put to sleep), or both. The legend shows the algorithm names as they appear in the respective works.

ward for port-based equipment (such as transponders). By using Power Aware Routing (PAR) instead of Shortest Path (SP) routing, additional savings can be achieved by concentrating backup paths and working paths on separate links, in order to be able to put the corresponding OLAs to sleep without being constrained by the presence of working paths. In Fig. 5 we show the power saving potential from applying both simple sleep approaches and PAR algorithms as reported by three works that allow to differentiate between these two approaches ([42], [43], [45]). Both Cavdar [42] and Jirattigalachote [43] estimate the combined power saving potential at around 30-40%. The work by Bao [45] reports savings which are more than 20 percentage points higher; however, this work also puts idle working path equipment to sleep¹⁰(which is not to be captured by our protection factor η_{pr}). While the work by Jirattigalachote [43] provides a detailed breakdown across the savings achieved by sleeping only (about 25% reduction, irrespective of the load¹¹), and PAR only (about 20% reduction at low load only), these results are not entirely consistent with the work by Cavdar [42] where the savings through sleeping are highly sensitive to the load, see Fig. 5.

In contrast to the previous three works that evaluated the power consumption at different loads, Coiro [44] considers a case study where the daily load varies according to a sinusoidal function. He uses only a PAR approach to power off the OLAs, and reports savings of 35%. While the above works only considered the WDM layer, the only work we found that considers both IP and WDM layer is by Musumeci et al. [46]. For the case of DPP, the work reports a 36–45% reduction in energy consumption; there is no breakdown over the impact of sleep mode and PAR. Inconsistent with the previous works, higher savings apply to higher traffic loads.

Summarized, the achievable savings through using sleep mode and PAR for protection equipment are roughly consistent around 30–40%, both for DPP and SPP; however, the savings profile is not always consistent among the various works.

Differentiated QoP — Ultimately, the customer-demanded level of reliability is a matter of cost. A reduction in protection power consumption could result from having cheaper SLAs that offer (slightly) less reliability with less-demanding customer requirements. This concept has also been branded as differentiated QoP. In [47] Lopez et al. calculate that by using differentiated QoP in a WDM transport network, on average savings of around 10% are possible with respect to DPP, regardless of whether it is a current fix-grid or an envisioned elastic (or flexible grid) network. The savings depend on the traffic load (higher traffic load leads to more savings), and the QoP levels required by the clients.

To assess the total protection factor power saving potential,

TABLE VII REDUCTION FACTOR FOR PROTECTION

		ing dualin	is nodes	S	×
	ANO	dr Sleeb	Diff	Subh	
Moderate Effort reduction					
IP	1	1	1	$1.00 \times$	
Opt. Swit	1	1.33	1	$1.33 \times$	
Transponder	1	1.33	1	$1.33 \times$	
OLA	1	1.33	1	$1.33 \times$	
Weighted total ^(a)				1.14 ×	(=-12%)
Best Effort reduction					
IP	1	1	1	$1.00 \times$	
Opt. Swit	1	1.66	1.11	$1.84 \times$	
Transponder	1	1.66	1.11	$1.84 \times$	
OLA	1	1.66	1.11	$1.84 \times$	
Weighted total ^(a)				1.29 ×	(=-23%)

^(a) See Section III-C for calculation details.

we again consider a Moderate Effort reduction factor and a Best Effort reduction factor. Even in a conservative estimation, employing sleep mode for protection devices seems to be the most promising solution, as it requires no additional management complexity for the operators. The above reported sleep mode savings of 25% correspond to a reduction factor of $1.33 \times$ (applicable to the WDM layer only). Thus, for our Moderate Effort reduction factor we assume that we have a reduction factor of $1.33 \times$ for the OXCs, transponders and OLAs. Taking the weight of the OXCs and transponders into account, this results in a **Moderate Effort reduction factor** of $1.14 \times$ (or 12% savings), see Table VII.

For our Best Effort scenario we assume that (a) that sleep modes and PAR exploit their full potential and can save 40% in the WDM layer (a reduction factor of 1.66), and (b) that differentiated QoP indeed leads to a 10% reduction (a reduction factor of 1.11) at the WDM layer. Note that the combined reduction potential in the WDM layer is $1.84\times$, which approaches the savings that could be achieved if we could do away with dualling in the WDM layer (or move to passive forms of dualling). As can be seen in Table VII this puts the **Best Effort reduction factor** for protection at $1.29\times$ (or 23% savings).

C. Amount of traffic T

The yearly growth in backbone traffic T is the major driver behind the continuous increase in power consumption in the backbone network. While the IP traffic growth in the backbone is not as high as it used to be—in the year 2000 traffic was growing at around 125% per year, but has slowed down to around 35% per year in 2010 [59]—it is still projected to grow at a rate that outpaces the improvements in energy efficiency of backbone telecom equipment, as shown in Fig. 1.

The traffic estimations in Fig. 1 are based on Cisco's twoyearly forecasts [60], [61], [59], [62]. The values for the year

¹⁰Bao [45] is (probably) able to do this because in their baseline scenario the network is not power-optimized at low loads (i.e. the power consumption of the baseline is identical for low and high load), whereas in [42], [43] the baseline power consumption at low load is lower than the baseline power consumption at high load.

¹¹The savings are *not* closer to 50% (as one might expect when considering that in a DPP scheme there is a backup path for each working path), because part of the OXC power consumption is considered static and thus not affected by the sleep mode.

 TABLE VIII

 TRAFFIC FACTOR RELATED ENERGY SAVINGS REPORTED IN PUBLICATIONS

App.	Source	Savings	Remarks	Justification
Data o	compression			
	Parker, 2011 [9]	50%	Probably overly optimistic	-
	Dong, 2012 [49]	45–55%	IP-over-WDM. MILP optimization on NSFNET. Upper range of savings for higher compression efficiencies. Study assumes compression ratios of 20:1 for video and 10:1 for images.	From Conclusion section in [49], which states 'optimizing data compression [] under the non-bypass approach [saves] up to 45% and 55% of the network power consumption'
	Kilper, 2012 [4]	negative	-	Section VI.B in [4] states that 'For the case of software- based compression using servers or PCs, uncompressed data transmission ($\sim 10^{-7}$ J/b for ten core hops) was shown to be more efficient than compressed data transmission after including the compression energy (> 10^{-6} J/b depending on the compression ratio [50])'
Cachi	ng			
	Lee, 2010 [51] Lee, 2011 [52]	10-60%	Evaluation over 20 content providers, con- siders complete network (i.e., CPE to core). Lower range of savings for DSL access tech- nology, upper range for GPON. Reason for considerable difference in savings between [51] and [52] is unclear.	From Fig. 4 in [51], <i>NonCCN</i> avg estimated at 2750, and <i>Core100%</i> avg estimated at 2250, gives 18%. Also states: '[<i>Core20% reduces</i>] energy consumption more than 15%.'. From Fig.4 and Fig 5 in [52] we estimate avg savings of <i>Core20%</i> and <i>Core100%</i> at 10% and 23% (DSL), and 25% and 60% (GPON) respectively.
	Osman, 2011 [53] Osman, 2013 [54] 8–37% IP-over-WDM. MILP NSFNET. Upper range of s demands and improved ca		IP-over-WDM. MILP optimization on NSFNET. Upper range of savings for higher demands and improved caching strategy.	Section 5.1 in [54] reports (consistent with [53]) that for the ' <i>fixed cache size</i> ' strategy daily-averaged savings ranging from 8% (low demands) to 30% (high demands). For the ' <i>variable cache size</i> ' strategy, the average savings range from 16% to 37%. (Results reported in the Con- clusion section are maximum instead of daily-averaged savings.)
	Chiaraviglio, 2011 [55] Chiaraviglio, 2012 [56]	45-65%	MILP optimization for 4 ISP backbone topologies. Reported savings probably cap- ture multiple effects beyond caching, as optimization algorithm also turns on/off servers and network nodes according to daily patterns and exploits server spare capacity, whereas non-energy optimized algorithm does not. Note that, in contrast to [54], low traffic conditions give higher savings, probably because of the above effects.	From Fig. 1.4 in [56], highest savings (low traffic) approx 65%, lowest savings (high traffic) approx 45%.
Modrzejewski, 2013 [57]		10%	Considers complete ISP network (i.e., access to core). Simulation on France Telecom network and Moroccan network.	Table II in [57] states 8.7% and 11.0%. Text also states: 'Energy savings of almost 9% and 11% are possible for the FT and the Moroccan scenarios, respectively'.
	Mandal, 2014 [58]	10–20%	ILP optimization for a hybrid peer-to-peer CDN, considers complete network (i.e., end- user, over access to core). Simulation on a US-wide IP backbone network and passive optical access network. Upper range for more popular content schemes.	Conclusion in [58] states that 'In some cases, our schemes can moderately reduce both server load as well as energy consumption (10%–20%)' However, the conclusion also reports an increase of energy consumption in other cases.

Acronyms: CCN: Content-centric Networking, CDN: Content Distribution Network, CPE: Customer Premises Equipment, DSL: Digital Subscriber Line, FT: France Telecom, GPON: Gigabit Passive Optical Network, ISP: Internet Service Provider, MILP: Mixed-Integer Linear Programming.

2010 and 2011 are taken from [59], while the most recent publication ([62], May 2013) provides a forecast for 2012 to 2017, and estimates the CAGR for the same period at 23%. We have extended this growth rate to the year 2020 (dotted line). As the annual projected growth rate is slowing down (at a rate of about 4 percentage points each year, based on [60], [61], [59], [62]), we have also plotted such a more likely, lower growth at 16% per year, from 2014 to 2020 (dotted line, labelled 'adjusted').

As can be seen in Fig. 1, even for the low traffic projection the backbone traffic increases 8-fold for the period 2010 to 2020. What are potential approaches to bring down the amount of traffic in the backbone network? The straightforward approach would be to set (reduced) quotas on the amount of traffic granted to each customer. However, this seems very unlikely from an economic perspective. Approaching the issue from a technical perspective, we will look at data compression and caching. An overview of the main works cited in the subsequent paragraphs is given in Table VIII.

Data compression — Data compression, or source $coding^{12}$, encodes information in such a way that it requires fewer bits than in the original representation. Parker et al. [9] estimate its potential for power reduction in photonic networking to be up to 50%. However, there are three pitfalls to be aware of. First, it seems unlikely that multimedia content (the bulk of global IP traffic [1]) can be further compressed, in or at the edge of the network. Multimedia content is already heavily compressed, for example audio mp3 compression, and video H.264 compression (used by default by, amongst others, YouTube and High-definition Television (HDTV) broadcasts). In [49], Dong et al. estimate the power saving potential through content compression at about 50%, however, they assume that video and images can be compressed 20 and 10 times respectively, which seems unlikely given that both media formats already widely employ compression at the application side. Second, compression and decompression at transmit and receiving side comes at a processing, and thus energy, cost. Kilper even states in [4] that uncompressed data transmission through the core is an order of magnitude more efficient than compressed data transmission. That said, the influence of the compression energy could drop with continuously more efficient Digital Signal Processings (DSPs) units [9], [66], and become negligible for content that is accessed by many users over a period of time. Third, any advances in compression techniques will probably be cancelled out by encoding more information and new modalities (such as stereoscopic view, or a higher dynamic range for audio or video) into bit streams

he historical evoluti

13

(i.e. Jevons paradox). To illustrate this, the historical evolution of video compression factors shows that for the latest main video compression standards, roughly each of them compressed twice as much [67]; the newest upcoming standard (HEVC), is again expected to continue this trend. However, despite this, the average bitrate per video stream has *not* consistently decreased, with a move from Standard-definition Television (SDTV) to HDTV (720 pixels), and then to HDTV (1080 pixels) [68]¹³.

Caching — Another technique to decrease the amount of traffic in backbone networks is the use of caching. The increase of media-rich internet content has lead to high bandwidth requirements for content served to multiple destinations. While the technique of caching is already well established-both at the client side, as well as between client and servers through intermediate proxy servers-Content Distribution Networks (CDNs) are the next logical step. A CDN is a large distributed set of servers deployed throughout the network, with content from the place of origin replicated to the other servers [69]. The main goal of CDNs is to increase availability and performance by serving requests from a server closer to where the request originates, but telecommunication service providers also deploy them to reduce the demand on their backbone. While caching content obviously consumes extra power, for the case study in [54] it is estimated that with optimal cache sizes a reduction of above 30% of the total power consumption (IP-over-WDM network and caching) can be achieved. The important issue to note here is that these savings are expressed taking into account both the reduced energy consumption in the network and the increased energy consumption in the caches. For the sake of argument we attribute all the savings to the network. Modrzejewski performs a similar study in [57]. For two realistic networks of national Internet Service Providers (ISPs) their algorithm to optimally decide where to cache content inside the ISP network predicts about 10% energy savings across the complete network, i.e. from access to core. The work by Lee et al. [52], that has coined the term 'content centric networking', considers network nodes that act as content caches themselves. They calculate potential savings over a network stretching from the customer side to a Tier 1 ISP core network. Evaluating for the top 20 content providers, they find that 10% to 60% can be saved over the complete network. However, as the large variation is very sensitive to the access network technology (DSL/GPON), this hints that the energy savings in the core will probably be rather limited. A very high saving potential is reported by Chiaraviglio et al. [56] (see Table VIII), but the evaluation captures savings effects from non-caching approaches as well, so these results are hard to interpret with respect to pure caching. Finally, the use of peer-to-peer caching could potentially reduce the caching power consumption overhead (and thus increase the overall savings), as such a scenario would exploit caching in end-user devices that are already on anyway, instead of integrating entirely new caches into the system as above. An

¹²Source coding should not be confused with *channel coding* and *network coding*. Channel coding in effect *adds* redundancy (instead of removing redundancy, as source coding does), to reduce the bit error rate in noisy communication channels. Network coding is a technique to attain the maximum possible information flow by algorithmically combining packets for transmission. There are studies available that look into network coding to improve the energy efficiency in wireless (ad-hoc) networks [63] and wired PONs [64], but almost none for backbone or transport networks. A single non-peer reviewed study [65] indicates energy savings in the order of 20% for two backbone networks, but states 0% savings for full mesh topologies. Given the limited amount of information available, we do not consider network coding at this point.

¹³Ultra high definition resolutions beyond HDTV are emerging, notably 4K UHD (3840x2160 pixels) and 8K UHD (7680x4320 pixels), but we did not find consistent and comparable data on the evolution of the average bitstream per video stream.

important issue in such a pure peer-to-peer caching scenario is that the energy consumption should not be migrated solely to the end-users, which would be very attractive for network operators but not necessarily reduce the overall (system-wide) energy consumption, as pointed out by Feldmann et al. [70]. A recent work exploring a hybrid solution, where a common CDN is combined with a peer-to-peer caching solution to obtain a more optimal overall system, is by Mandal et al. [58]. The system-wide energy savings (network+caches) are reported as 10–20% in some cases; however other cases are reported as requiring an equal amount or even more energy. Furthermore, a breakdown of the power consumption savings and increase over the different network sections is not given.

Following the above findings, practical approaches to reduce the amount of traffic in the core network seem to be limited. Therefore, for the **Moderate Effort reduction factor** we assume no savings, i.e. a reduction factor of $1.0\times$. For our Best Effort scenario, we assume that indeed a reduction of about 10% of energy consumption can be achieved both through compression and energy-optimal content caching. This means that we have a **Best Effort reduction factor** of $1.23\times$ (or 19% savings). Both values are summarized in Table IX.

 TABLE IX

 Reduction factor for amount of traffic



D. Power rating P/C (part I)

The power rating factor P/C expresses equipment power consumption as the power per unit capacity, e.g. 5 W/Gbps. While we defined the denominator as the equipment capacity, we will relax this constraint and also include in this section any approaches that improve the power consumption under varying load. The baseline values that we assumed for the various components are listed in Table II.

Fig. 6(a) shows how the power consumption of network equipment typically scales with varying load. Note that the maximum load is equal to the capacity of the equipment. Several works have identified that for network switching equipment the power consumption when idle is around 90% of that at maximum load. Chabarek et al. [74] measured the power consumption of two Cisco routers in different configurations and under various loads. They observed that while the configuration (chassis and line cards used) significantly influences the power consumption of the router, the load has limited impact. These observations are confirmed by Mahadevan et al. [75], and in our earlier work [20]. In addition, in [20] we showed that this is also the case for OLAs. While to our knowledge there is no public data available for other optical equipment, such as WDM terminals, transponders, or OXCs, it seems to be accepted as a fact that this trend holds there as well.

There are two general categories of approaches to reduce the power consumption shown in Fig. 6(a), and thus the power rating factor.

- The first category, shown in Fig. 6(b), focuses on reducing the power associated with the maximum load, i.e., the device's capacity. In this category, we will consider the *inherent component energy efficiency improvements* that are observed for communication equipment year after year, and *improved chassis utilization* through better chassis and shelf filling.
- The second category, shown in Fig. 6(c), focuses on making the equipment more power proportional, i.e. the power scales better with the actual load. The typical techniques to do so involve putting (sub)components or systems to *sleep*. This can be applied on *short time scales* (in the order of packet-level transmission) and on *long time scales* (in the order of hours and days).

For ease of reading, we will discuss the first category (more efficient components, and improved chassis utilization) in this section. The second category (sleep mode approaches) will be moved to, and discussed in, the subsequent section (Section IV-E). An overview of the main works cited in this section is given in Table X.

More efficient components — Telecommunication equipment becomes more power efficient each year, largely driven by more energy-efficient CMOS technology. However, as already observed by Kilper [1] and Tucker [76] the rate of improvement has been slowing down. While up to 2006 the energy-perbit reduction (i.e., J/bit or W/Gbps) in routers was around 20% per year [71], it has slowed down to around 13% per year more recently [48], see also Fig. 7. In [1], it has been argued that this efficiency improvement might further slow down as a result of practical limitations inherent to CMOS transistor design. Koomey et al. [66], on the other hand, argue that the efficiency improvement trend for electronic processing equipment will continue through clean slate design, similar to what was observed a few decades ago with the transition from vacuum tubes to discrete transistors and subsequently to microprocessors. Note that the reported efficiency trends in Table X and Fig. 7 apply mainly to (opto)electronic equipment.



Fig. 6. Power scaling with load and capacity. (a) For typical network equipment the power consumption when idle is still close to the that at maximum load. (b) More energy-efficient equipment consumes less energy at maximum load. (c) Perfectly power proportional equipment scales linearly with the load.

 TABLE X

 Power rating factor related energy savings reported in publications (part I)

App.	Source	Savings	Remarks	Justification					
More	More efficient components (annual energy-per-bit reductions)								
Neilson, 2006 [71] 20% p.a.		20% p.a.	Trend up to 2006 for high-capacity routers. Note that our trend line in Fig. 7 gives 29% p.a., whereas Neilson writes 20% p.a.	Caption of Fig. 4 in [71].					
	Han, 2010 [72]	15% p.a.	Trend from 1992 to 2008 across 3 generations of Fujitsu optical transport platforms (ADMs \rightarrow MSPPs \rightarrow POTPs)	Section II in [72]: 'power consumption for each transported gigabit per second has decreased over the past 16 years, from 84 W/Gb/s for an ADM to 6.7 W/Gb/s for a POTP'					
	Tamm, 2010 [48]	13% p.a.	Trend and projection for 2005 to 2020. Applies to routers, packet switches, SDH-XCs, and OTN-XCs.	Fig. 7 in [48].					
	Lange, 2011 [2]	14% p.a.	Trend and projection for 2002 to 2020 based on routers from 3 different vendors.	Fig. 10 in [2].					
Impro	oved chassis utilization								
ľ	Ceuppens, 2009 [73]	55%	Difference in power rating of specific Juniper router chassis filled with 1 slot compared to all slots filled.	Slide 6 in [73] for MX960 Juniper router: ECR with 1 line card is 18 W/Gbps, with 11 line cards is 8 W/Gbps.					
	Van Heddeghem, 2012 [20]	46% (IP eq) 46% (OLA) 9% (WDM)	Savings from suboptimal chassis/shelf filling over usage lifetime for IP equipment, OLAs and WDM terminals, respectively.	Fig. 6 in [20].					

Acronyms: ADM: Add/Drop Multiplexer, ECR: Energy Consumption Rating, MSPP: Multiservice Provisioning Platform, OTN: Optical Transport Networking, POTP: Packet Optical Transport Platform, SDH: Synchronous Digital Hierarchy, XC: Cross-Connect.

Components that are dominantly optical (such as OXCs and OLAs), improve at a much slower rate [1], although we did not find credible values for these two components. Given the above observations, a 13% per year reduction in energy-perbit for electronic and optoelectronic backbone equipment (IP routers and transponders) in the time frame 2010 to 2020 seems reasonable. For dominantly optical components (OXCs and OLAs) we assume no yearly improvements. Note that a 13% per year reduction in energy-per bit (i.e., J/bit or W/Gbps) corresponds to a 15% improvement in energy efficiency (i.e., bit/J or Gbps/W), since 1/(1 - 0.13) = 1.15. These are the values reported in Table XII.

For our complete backbone communication stack, the resulting yearly reduction in energy-per-bit (W/Gbps) then becomes 11%, or a corresponding 13% efficiency improvement (Gbps/W) which is the improvement rate plotted in Fig. 1. It is important to point out that this is an idealized improvement rate, and would only be achieved in practice if all equipment in the network is continuously replaced by the latest (i.e., most efficient) generations, which is obviously unrealistic¹⁴.

At this point, it should be noted that the yearly energy efficiency improvement of network equipment is the only power saving approach in this study that is time-dependent. All other approaches, both those already discussed and those yet to be discussed, are *static* (once-only) approaches; their full energy reduction potential can only be applied once, and not reused. Because of the time-dependency of more efficient components, we will treat the impact of a 11% p.a. reduction in energy-per-bit (= 13% efficiency improvement) separately in Section V-A2.



Fig. 7. Evolution of power rating values of various telecom equipment as given by Neilson [71], Han [72], Tamm [48], and Lange [2]. Exponential trend lines have been added. Note that some data points are actual values, while others are projections; the publication year might give an indication.

Improved chassis utilization — While Chabarek et al. already reported in [74] that 'from a power-aware perspective, it is best to [...] maximize the number of line cards per chassis.', to our knowledge the work by Ceuppens [73] is one of the first public works that considers the impact of slot filling levels on the *power rating* value of an IP router. Based on measurements on an MX960 Juniper router, he finds that 'chassis utilization below 30% significantly affects [the power rating]', i.e. the chassis with only one line card is power rated at a value twice as high as the one filled with all line cards. In [20], we have considered this issue more in detail, and looked into the impact of equipment filling levels on the power rating value of both IP equipment and WDM equipment. Equipment deployed in the field is not always optimally filled; but instead often starts off with an almost empty chassis which, over time, is filled with more and more line and control cards. As a result, power rating values will only approach their optimal (lowest) value towards

¹⁴In this context, the work by Parker et al.[77] provides an interesting study on energy-efficient upgrade paths for new router generations through mastslave configurations; however, the work applies to edge routers instead of backbone routers.

the end of the equipment's life, when the chassis overhead is shared by the maximum number of functional components. The results in [20] indicate that the optimal power rating achieved when the rack is at maximum capacity, i.e., the power rating value typically assumed in power models-in some cases needs to be corrected by close to a factor 2 (i.e, twice as worse). Particularly, the study found that IP routers over their lifetime are a factor of 1.85 more inefficient than at optimal (i.e., fully filled) capacity. Otherwise said, if the power rating of IP routers would scale better with the filling level, they would be 1.85 times more efficient (or save 46%) than is currently the case. As this value is sensitive to the lifetime assumptions, we assume a slightly lower **Best Effort reduction factor** at the IP layer of $1.5 \times$ (or 33% savings). Based on the same work [20], we assume for OLA equipment a reduction factor of $1.5 \times$ (or 33% savings), and a reduction factor of $1.1 \times$ (or 10% savings) for WDM terminals¹⁵. Note that the WDM terminals in [20] include the transponder power consumption, so we apply the same factor there. For the Moderate Effort reduction factors for the above three layers, we assume no savings, i.e. a reduction factor of $1.0 \times$. For a summary, see Table XII in the next section.

E. Power rating P/C (part II)

This section discusses the sleep mode approaches (on a daily time scale, and on short time scale) to improve the power rating factor P/C. It is a continuation of the previous section, but has split up for ease of reading. An overview of the main works cited in this section is given in Table XI.

Sleep modes on a daily time scale — A popular research topic to improve energy efficiency is the usage of sleep modes, and examples of such works are ample (see the surveys [6], [7]). We already discussed it in the context of protection as well (Section IV-B). The general idea is based on the fact that even in core networks communication equipment is not always working at maximum load. However, power consumption of communication equipment remains almost independent of the actual load ([74], [75], [20]), as illustrated in Fig. 6(a). As such, making the equipment or system more power proportional by shutting down (sub)components when it is not in use could lead to substantial overall savings. Before expanding on the two different sleep mode approaches to improve power proportionality, it is instructive to look at the daily traffic variation. As shown by Lange et al. in [78], the traffic volume in communication networks varies considerably over time. For aggregated traffic (as opposed to service-specific traffic), the variations range from peak values typically in the evening, to off-peak values as low as 25% in the morning (see the inset in Fig. 8). Other actual traffic matrices showing this daily variation can be found in [81] (France Telecom network), [93] (anonymized), and [94] (showing the interesting fact that European network traffic reaches lower off-peak values compared to U.S. traffic). The aggregated traffic in the core can be approximated by





100%

Fig. 8. Theoretical achievable sleep mode savings S for a sinusoidal traffic variation over the day. The current off-peak traffic is estimated by [78] at $\alpha = 0.25$, which corresponds to maximum theoretical savings of 38%. The sine curve has been phase shifted to approximate the traffic curve from [81] with a minimum around 5am.

a sinusoidal curve, and for current networks the minimum traffic value is in the range of 25% of the peak traffic volume [78]. Note that from a theoretical perspective, if the load exhibits a perfectly sinusoidal variation over the day, with its minimum being zero during off-peak times, the savings will be upper bounded to 50%. More generally, if the offpeak load is α instead of zero, then the theoretical savings are $S(\alpha) = 0.5 \times (1 - \alpha)$, see Fig. 8. With the off-peak load being estimated in current core networks at α =25%, the theoretical upper bound to the savings are thus 38%. For more realistic stepwise adaptivity to the load, the savings are up to 19% [78]. Similar to what we saw for exploiting sleep modes for protection equipment (Section IV-B), we consider a relatively straightforward *simple turn off* strategy, and a more elaborate PAR strategy to maximize the equipment that can be put to sleep; both discussed next.

The most simple approach to exploit this theoretical potential is to simple turn off equipment when it is idle. Visualized on Fig. 6(a) this would mean that the power of the relevant equipment at idle drops from P_{min} to P_0 . If this can be done stepwise for a number of subcomponents (such as the line cards in an IP router), then we can gradually approach the power proportionality shown in Fig. 6(c). The premise for being able to simply turn off links, is that one logical (IP) link is actually a 'bundled link'¹⁶ that consists of multiple aggregated sublinks. For example, a 40 Gbps bundled link may be realized through four 10 Gbps interfaces or sublinks. The study by Fisher et al. [83] claims that 'in today's backbone networks, a vast majority of links would be bundled, with bundles consisting of two to approximately twenty cables [i.e., sublinks], a majority between the two extremes'. It is argued in the same work that the drivers behind link bundles are both (a) capacity requirements exceeding the fastest available link technology, and (b) the fact that capacity upgrades are often realized by adding new links alongside existing ones (rather than replacing the existing equipment with a higher-capacity link). Doverspike

¹⁵In contrast with the IP equipment, for OLAs and WDM terminals we don't downward adjust the savings reported by [20] in Table X, as they are not so sensitive to a lifetime parameter as is the case with the IP equipment.

¹⁶Link bundling is also referred to under various other umbrella terms such as link aggregation, link bonding, link teaming and port trunking. The IEEE 802.1AX-2008 standard uses the term 'link aggregation'.

 TABLE XI

 Power rating factor related energy savings reported in publications (part II)

App.	Source	Savings	Remarks	Justification			
Sleep	Sleep modes – daily time scale						
	General Lange, 2010 [78]	19–38%	Models theoretical upper bound of savings based on three daily traffic profiles (layer agnostic). Higher savings apply to perfect power proportional networks, lower savings for more realistic stepwise adaptivity.	Section 4 in [78], savings for aggregated (sinusoidal) traffic.			
	Simple turn off Liu, 2011 [79]	86% (65%) ^(a)	IP router ports only. Savings for turning off sublinks at low utilizations in a daily traffic profile, on synthetic Internet2- based network. Reported savings also capture (probably large) effect of exploiting overprovisioning.	Section VII in [79].			
	Idzikowski, 2011 [80], (Idzikowski, 2013 [81])	2–35% (2–26%) ^(a)	IP line cards only. MILP optimization for three backbone topologies. Savings for turning off sublinks of link bundles ('parallel line cards') at low utilizations. Upper range of savings for higher demands and more gravitational traffic models.	In [80], comparing estimated daily average of ' <i>FUFL</i> ' with ' <i>Static Base Network</i> ' in Figures 6(a) and (c), and Figures 9(a) through (d). Table 4 in [81] reports 15% for MUELL=0.5.			
	Power Aware Routing Chabarek, 2008 [74]	2–65%	Savings <i>probably</i> for router chassis + line card. Savings for rerouting/grooming at low utilizations in a daily traffic profile. High savings for single-port line cards and dense networks.	Table IV in [74].			
	Restrepo, 2009 [17]	10%	IP layer only. MILP optimization in 50-node core network for EPAR based on five distinct energy profiles. Exploits traffic differences among different nodes.	Fig 6 in [17], for <i>On-Off</i> and <i>Log100</i> (approaches profile of current routers)			
	Zhang, 2010 [82]	30%	IP layer only. MILP optimization on NSFNET case study. Savings for turning off chassis and line cards at low utiliza- tions in a daily traffic profile.	Section 3 in [82], daily averaged power savings: 'The power savings varies from 0% to 68.7%, with an [daily] average value of 29.8%'. Study considers router chassis and line cards.			
	Fisher, 2010 [83]	40–80% (30–60%) ^(a)	Probably only applicable to IP router line cards (power model not given). ILP optimization and heuristics. Savings for turning off sublinks in link bundles at low utilizations in a daily traffic profile. High savings apply to larger bundle sizes. Reported savings also capture (probably large) effect of exploiting overprovisioning.	From Fig 2 (Abilene) and 3 (two syn- thetic topologies) in [83].			
	Idzikowski, 2013 [81]	39–45% (29–34%) ^(a)	IP line cards only. Survey and benchmark of 6 different PAR approaches (covering the works [84], [85], [86], [87], [88]) exploiting daily traffic variations on France Telecom reference scenario.	Table 4 in [81] for MUELL=0.5. Higher values of MUELL seem undesirable for network operators.			
Sleep modes – short time scale							
Ĩ	Nedevschi, 2008 [89]	40–60% (30–45%) ^(a)	IP line cards only. Savings for inter-packet sleeping at low utilizations using a buffer-and-burst approach. Packet-level simulation on two realistic topologies and traffic workloads (Abilene and Intel), assuming at packet-level a mix of Pareto- flows and constant bit-rate traffic . Upper range of savings for low utilizations.	Figures 14(b) and (d) in [89] are representative for current equipment with high idle power. Savings from 60% (utilization=10%) to 40% (util.=30%)			
	Reviriego, 2010 [90] Reviriego, 2011 [91]	-	Evaluates the feasibility of Energy-Efficient Ethernet (EEE) on optical high-speed links.	-			

Acronyms: EPAR: Energy Profile Aware Routing, FUFL: Fixed Upper Fixed Lower, MILP: Mixed-Integer Linear Programming, MUELL: Maximum Utilization of Each Logical Link, PAR: Power Aware Routing.

^(a) Estimated savings at IP layer, using the finding from [92] that line cards represent about 75% of the IP power consumption.

et al. [95] mention a third driver for link bundling, which is resilience and consequently network stability; if one of the component links fails the bundled link remains up and a failure-driven topology update is not required. Unfortunately, Fisher et al. [83] do not provide actual data (such as link bundle counts for real operators) to ground their—otherwise plausible—claims. Two works that consider a simple turn off strategy of sublinks (without changing the logical topology) are by Liu and Ramamurthy [79] and Idzikowski et al. [80]. Liu and Ramamurthy [79] report 86% IP port savings by putting sublinks to sleep following daily traffic variations. As line cards represent about 75% of the IP router power consumption [92], this would mean that in the IP layer savings around 65% are possible. The reason for this very high savings percentage (much higher than our indicative, theoretical upper value of 38%, as derived earlier) is that they exploit overprovisioning in the network; they report for a specific link that the 'utilization is commonly less than 1%'. We should be careful not to assign the potential saving effects of exploiting overprovisioning, as this is often installed with good reason (see further). The study by Idzikowski et al. [80] uses a similar technique in its 'FUFL' algorithm to turn off sublinks (referred to as 'parallel line cards') at low utilizations, resulting in savings ranging from 2% to 26% (IP layer savings) depending on the demand size and gravity of the demand model. This study *does* consider keeping the overprovisioning factor (captured by the Maximum Utilization of Each Logical Link (MUELL) parameter) identical in both the baseline scenario as well as

the sleep scenario. So these estimations are more useful for our purpose. The practical advantage of the simple turn off strategy lays in the fact that all decisions can be taken locally and require no logical topology changes, which is especially attractive for network operators where network stability is critical. A significant drawback includes the fact that the actual saving potential relies heavily on the amount of bundled links being prevalent in current and future core networks¹⁷.

A more elaborate approach, which we will refer to as Power Aware Routing $(PAR)^{18}$, is to not wait until a line card is completely idle, but instead pro-actively reroute traffic so that traffic of lightly loaded links is moved to links with spare capacity. In practice this means that traffic is groomed. This can be done periodically (e.g. multiple times per day, or even per hour) in response to the demands changing throughout the day, resulting in a reconfiguration of the logical topology. The obvious gain of the extra effort involved in PAR is to increase the amount of energy savings that are possible in response to these traffic variations. An important aspect here is the link overprovisioning, also called the maximum link load. Network operators overprovision the capacity of backbone links by a factor of 2 or even more ([1], [97]) to account for peak-to mean traffic variations, unexpected traffic spikes (e.g. in response to a major news event), and future traffic growth (see Section IV-G). Any study evaluating the energy saving potential of PARlike approaches should make sure that the overprovisioning factor is kept identical in both the baseline scenario and the power optimized PAR scenario to allow for a fair comparison. As was shown in Figure 17 of [85] and in Table 4 of [81] (with the overprovisioning factor labelled MUELL), increasing the overprovisioning factor only in the optimized scenario results in significantly higher energy savings (more than 20 percentage points), with the hidden cost being a reduction in Quality of Service (QoS). The issue in some studies is that this influence is not always properly accounted for or clearly indicated¹⁹, which might lead a reader to interpret results overly optimistic. In the context of sleep mode approaches to save energy (both simple turn off, and PAR), the work by Idzikowski et al. [80] is specifically interesting as it both (a) has a good survey on earlier related work highlighting some important open issues in these works, and (b) investigates in a consistent manner the sleep mode saving potential taking into account consistent overprovisioning. Furthermore, some of the algorithms proposed in [80] are evaluated on a testbed by Tego et al. [98], which so far has only rarely been done for research in this area. For the saving potential of PAR approaches we will however refer to the more recent work by the same author [81], as it evaluates different PAR algorithms

(notably from [84], [85], [86], [87], [88]) across a consistent scenario (topology and traffic matrix). The reported savings for the different algorithms are in a very narrow range (29% to 34%, IP layer savings, see Table XI). The other PAR-related works we have listed in Table XI ([74], [17], [82]) report savings consistent with that finding. A notable exception is Fisher et al. [83] who report much higher savings (40% to 80%, probably applicable to IP router line cards, which would correspond to IP layer savings around 30% to 60%). This is most likely because the work also exploits overprovisioning. Summing up, the savings attributable to sleep modes that exploit daily traffic variations, the 'simple turn off' strategy seems attractive both from the real savings it can offer and the limited added technical and operational complexity required. Thus, for our Moderate Effort reduction factor at the IP layer we assume $1.15 \times$ (or 13% savings), as an average of the IP layer savings estimated by [80]. With the more operationally challenging PAR we assume for the Best Effort reduction **factor** the upper bound from [81], i.e. $1.5 \times$ (or 34% savings) not only at the IP layer but also at the other layers. For a summary, see Table XII.

Sleep modes on a short time scale — While the approaches discussed in the previous section exploit traffic variations on a daily time scale, there is additional energy saving potential on a packet-level time scale which can be partially independent from the previous approach. In packet switched networks, packet arrival rates can be highly non-uniform. In between packet arrivals, it is possible to off turn (sub)components. However, as the transition from operating to sleep state consumes both energy and time (which might result in an increase in delay), there is a trade-off to the extent where this is possible. While there are several works on applying this approach in Ethernet LAN switches (e.g., Gupta [99], quoting savings in the range of 30% to 80%), we are aware of only the work by Nedevschi et al. [89] that considers it in the context of core networks, as shown in Table XI. They base their evaluations on packetlevel simulations with real-world network topologies and traffic workloads. Their findings for the baseline of equipment where idle-mode power is almost comparable to active-mode power (as is the case for current telecom equipment) are savings in the range of 30% to 45% (IP layer savings, see Table XI). The above described approach is very similar to Energy-Efficient Ethernet (EEE), standardized as IEEE 802.3az, but which targets 100 Mbps to 10Gbps Ethernet copper-based physical layer devices. However, Reviriego et al. [100] found that already for link speeds of 1 Gbps and higher the proposed transition times (i.e., sleep and wakeup times, both in the order of microseconds²⁰) are too high to achieve reasonable power proportionality. In addition, traffic patterns play a key role in the achievable savings. For optical transceivers, the transition times are even much larger (in the order of 1–2 milliseconds) than for copper-based devices due to a more complex circuitry to stabilize the channel, although recent work suggests that they can be dramatically reduced [90]. For high speeds optical

¹⁷It is interesting to point out that we identified the same high sensitivity to the ratio of the required link capacity over the available interface capacity in [96] (referred to as demand/linerate ratio) in the context of optical bypass.

¹⁸Note that similar approaches have been referred to in other works as Energy Aware Routing (EAR), Energy Aware Adaptive Routing (EAAR) or Green Routing.

¹⁹For example, in [86] it might not be instantly clear that the cost savings (which can act as a proxy for energy savings) listed in Table VII are for different overprovisioning values (maximum lightpath utilization $\delta = 1.0$ for LFA and GA, and maximum utilization of last lightpath on a logical link $\psi = 0.9$ for EWA) than the reference scenario (overprovisioning factor $\gamma = 0.5$).

 $^{^{20}}$ For 10GBASE-T, the EEE wake and sleep transition times defined by IEEE 802.3az standard are 4.48 μ s and 2.88 μ s respectively, while the transmission time for a 1500 byte frame is 1.2 μ s [100].

 TABLE XII

 Reduction factor for the power rating factor

		compon	ans ⁶⁰ uil	Ration de	ally shot	, ,
	More	h. mpr.	Sleep	nic Sleep	no Sibitote	
Moderate Effort red						
IP	(1.15)	1	1.15	1.25	$1.44 \times$	
Opt. Swit	-	1	1	1	$1.00 \times$	
Transponder	(1.15)	1	1	1	$1.00 \times$	
OLA	-	1	1	1	$1.00 \times$	
Weighted total ^(a)	(1.13)				1.18 ×	(=-15%)
Best Effort reduction						
IP	(1.15)	1.50	1.50	1.40	3.15×	
Opt. Swit	-	1.10	1.50	1	$1.65 \times$	
Transponder	(1.15)	1.10	1.50	1.40	$2.31 \times$	
OLA	-	1.50	1.50	1	$2.25 \times$	
Weighted total ^(a)	(1.13)				2.62 ×	(=-62%)
(a) Son Sontion III C f	or coloul	ation day	toile			

^(a) See Section III-C for calculation details.

^(b) The effect of more efficient components resulting from yearly efficiency improvements is *not* included in the weighted total in this table. See Section V-A2 for impact of yearly efficiency improvements.

links (40 Gbps and above), a promising alternative approach to circumvent the above issues with the wake/sleep transitions might be the use of a multilane architecture. That is, realizing parallel lower-rate lanes instead of one single high-rate lane; e.g. 4×10 Gbps to realize a single 40 Gbps link, not unlike the 'link bundling' concept discussed before. Depending on the traffic load, a number of lanes can be powered down and thus save energy. This has been studied in the context of energy efficiency by Revigiero et al. [91]; their results show indeed an improved power-load profile. So, while there are certainly some issues to resolve with respect to sleep modes at a packet-time scale, there do seem to be feasible solutions on the mid-term to realize important savings here. As the interdependency with sleep mode at daily time scale is not completely clear-cut, we assume for the Moderate Effort reduction factor $1.25 \times$ (or 20% savings) at the IP layer only, being a slightly lower value than the lower bound reported by Nedevschi et al. [89]. Similarly, for the Best Effort reduction factor we assume $1.4 \times$ (or 30% savings) for the IP layer and transponders (it seems unlikely that OXCs and OLAs can also benefit from this in the mid-term).

Applying the proper weights to the values in Table XII, the overall result for the power rating P/C factor is a **Moderate Effort reduction factor** of $1.18 \times$ (or 15% savings), and a **Best Effort reduction factor** of $2.62 \times$ (or 62% savings), as shown in the same table.

F. Hop count H

The layer hop count H represents the average number of hops between processing elements in the respective layer, for example IP nodes in the IP layer. For a given network topology the hop count will depend on several aspects, such as the routing algorithm and link weights. However, a good ballpark number for H in a backbone network is 3–4 hops [1], [16], [21].

We consider the technique of optical bypass to reduce the hop count. An overview of the main works cited in the subsequent paragraph is given in Table XIII.

Optical bypass — A well-known technique to reduce the hop count H in the IP layer is to optically bypass IP routers, also known as IP offloading. The idea is that traffic not intended for the IP node remains in the optical layer and thus bypasses the IP router. The lightpath is switched, using OXCs, from an incoming fiber link directly on the appropriate outgoing fiber link. This allows us to reduce the capacity of the router and the associated power consumption. Optical bypass is possible at single-wavelength granularity, or on waveband granularity (requiring fewer ports in the OXC since multiple wavelengths are switched at the same time). The first thorough study to investigate optical bypass in the context of energy efficiency was by Shen et al. [101]. Next to an optimal upper bound on the savings, they evaluated two heuristics. While their 'direct bypass' strategy achieves savings around 5-45% (low-high demand), the 'multi-hop bypass' strategy (an intermediate, hybrid solution) was able to increase these savings with around 10 percentage points at lower demands.

From an energy (and also cost) perspective, optical bypass requires adequately filled optical channels, an issue not tackled in [101]. We have shown in [96] that if the lightpath filling ratio²¹ is roughly below 50%, optical bypass consumes more energy than IP switching and grooming. Furthermore, in the same work [96] we've identified that another important parameter is the mesh degree of the network; full-mesh topologies benefit less from optical bypass.

As for small, full mesh core networks that serve traffic through high capacity interfaces (i.e., with demands relatively being small) the savings potential from employing optical bypass is limited (or even negative), we assume for the Moderate Effort reduction factor no savings, i.e. a reduction factor of $1.0\times$. On the other hand, if the average hop count in core networks is indeed around 3-4 hops, then there is considerable saving potential at the IP layer and transponders. Therefore, for the Best Effort scenario we have applied a factor 3 to the IP layer and transponders, as shown in Table XIV. This results in a **Best Effort reduction factor** of $2.03 \times$ (or 51%) savings), which is consistent with the results found in the above cited work. We should note that some PAR approaches from Section IV-E might actually increase the hop count slightly, and therefore reduce the combined saving potential, which is one reason why we used the lower bound of hop count reductions at 3 instead of 4. Both reduction factors are summarized in Table XIV.

 $^{^{21}}$ Note that lightpath filling relates to the concept of link bundles discussed in Section IV-E. For example, a link bundle consisting of 3 links would imply a lightpath filling ratio of 300%. However, the link bundle term cannot capture lightpath filling ratios < 100%.

 TABLE XIII

 HOP COUNT RELATED ENERGY SAVINGS REPORTED IN PUBLICATIONS

App.	Source	Savings	Remarks	Justification
Optical Bypass				
-	Shen, 2009 [101]	15–45%	IP-over-WDM. MILP optimization and heuristics on three backbone networks (a six-node eight-link network, NSFNET and USNET). Upper range of savings for higher demands, and networks with more nodes.	Fig. 5 in [101], for the ' <i>multi-hop bypass</i> ' heuristic, which outperforms the more simple ' <i>direct bypass</i> ' heuristic.
	Van Heddeghem, 2013 [96]	0–75%	IP-over-WDM. Shortest path routing across a variety of scenarios. Upper range of savings for sparse networks and high demand/linerate ratios.	Fig. 3–6 in [96].

Acronyms: MILP: Mixed-Integer Linear Programming.

TABLE XIV REDUCTION FACTOR FOR HOP COUNT

	Optice	LI DYPass tote	>
Moderate Effort reduction	1	1.00×	(=-0%)
Best Effort reduction			
IP	3.00	$3.00 \times$	
Opt. Swit	1	$1.00 \times$	
Transponder	3.00	$3.00 \times$	
OLA	1	$1.00 \times$	
Weighted total ^(a)		2.03 ×	(=-51%)

^(a) See Section III-C for calculation details.

G. Overprovisioning factor η_{op}

 TABLE XV

 Overprovisioning factors reported in publications

Source	Overpr.	Remarks
Zhang-Shen, 2004 [97]	10×	Given as ' most networks today are enormously overprovisioned, with typical utilizations around 10%.'
Huang, 2005 [102]	\geq 5×	Given as 'to accommodate the relatively large traffic fluctuations caused by link failures, values of $\beta \approx 5$ or even higher are not uncommon in large IP backbones'
Fisher, 2010 [83]	2–3×	Given as 'The capacity of backbone net- works is overprovisioned []. The aver- age link utilization in backbone networks of large Internet service providers is esti- mated to be around 30–40%.'
Kilper, 2011 [1]	$2 \times$	Does <i>not</i> include overprovisioning for protection.
Kilper, 2012 [4]	2–4×	Given as 'In the core network, the traffic is more uniform and the overprovisioning ratio can be as low as a factor of 2– 4 times', however this factor includes overprovisioning for protection as the au- thor states 'overprovisioning is used to traffic bursts, provide spare capacity for use in the event of a failure (protection), and allow for mean traffic growth over a period of time'

We have already mentioned the practice of overprovisioning before. Network operators typically overprovision the capacity of their backbone to account for peak-to-mean traffic variations, unexpected traffic spikes (e.g. in response to a major news event), failures and future traffic growth. We have separated the overprovisioning related to failures (i.e., protection) in the protection factor η_{pr} , and dealt with it in Section IV-B. The reason we did so, is because there are numerous works that investigate specifically the energy saving potential of overprovisioning for protection, allowing a more accurate assessment of the power reduction potential.

In Table XV we list the backbone overprovisioning factors as reported by a number of sources. From the values given, it is clear that there is a rather wide range of estimates on the actual overprovisioning factor. However, it is safe to assume that most of these factors should be halved, as they include the overprovisioning for protection; the work by Kilper et al. [1] is a notable exception. After doing this, the lower range of overprovisioning is a factor of 2.

The power consumption associated with this overprovisioned capacity can be reduced with the sleep mode techniques we discussed earlier in Section IV-E. For the Moderate Effort scenario we assume that at least a factor 2 can be saved at the IP layer, which results in a Moderate Effort reduction factor of $1.34 \times$ (or 25% savings). For the Best Effort scenario, we assume that a factor of 2 can be applied to all components in the network, resulting in a Best Effort reduction factor of $2 \times$ (or 50% savings). This is summarized in Table XVI.

TABLE XVI REDUCTION FACTOR FOR OVERPROVISIONING

	Steel	modes subtote	>
Moderate Effort reduction			
IP	2.00	$2.00 \times$	
Opt. Swit	1	$1.00 \times$	
Transponder	1	$1.00 \times$	
OLA	1	$1.00 \times$	
Weighted total ^(a)		1.34×	(=-25%)
Best Effort reduction	2.00	2.00 ×	(=-50%)
	1.1.1		

^(a) See Section III-C for calculation details.

V. TOTAL POWER SAVING POTENTIAL AND DISCUSSION

In this section we combine the power saving potential of the different approaches into a single number for the Moderate Effort and Best Effort scenario. We perform a sensitivity analysis on the two parameters in our model from Section III that could not be abstracted away, i.e., the average hop count H and the average link length. We compare our findings with the savings potential estimated by Greentouch's Green Meter [10]. Finally, we briefly assess how legacy networks (in contrast to our IP-over-WDM network) would impact our findings.

A. Total power saving potential

For discussing the impact of the reduction factors, we will distinguish between the static (once-only) approaches, and the time-dependent yearly efficiency improvements.

1) Static approaches only: In Fig. 9(a) we show the reduction factors for each of the static approaches, both in the Moderate Effort scenario (transparent bars) and the Best Effort scenario (filled bars)²². The time-dependent annual efficiency improvements are not included, but will be accounted for in the next section. The approaches with a high potential stand out clearly, i.e. reducing the overhead factor, the employment of optical bypass, and sleep modes for the overprovisioned capacity. However, to assess the realistic impact of each of the approaches, it makes sense to look at those that perform relatively well in *both* the Moderate Effort and the Best Effort scenario. In this light, a reduction of power associated with the overhead factor and the overprovisioning, again, seems a promising and realistic direction. Another approach that performs well in both scenarios are sleep modes in general.

If we combine all of the static approaches depicted in Fig. 9(a), not yet including the time-dependent annual efficiency improvement, the total reduction potential in the **Moderate Effort scenario** is $2.3 \times$ (or 57% savings), and in the **Best Effort scenario** $31 \times$ (or 97% savings). Note that multiplying each of the reduction factors from Fig. 9(a) is a good approximation to get the combined reduction potential²³.

2) Static and time-dependent approaches: However, we have not yet accounted for the 11% yearly reductions in energy-per-bit (or 13% improvements in energy efficiency, which makes for a yearly reduction factor of 1.13), see Section IV-D. In Fig. 9(b) we plot for the time frame 2010 to 2020 the combined potential of both the static reduction factors (discussed above) and the impact of the expected yearly efficiency improvements. In both scenarios, the annual efficiency improvement is able to increase the static reduction potential with a little more²⁴ than a factor 2, resulting in a **Moderate Effort** reduction factor of $6.0 \times$ (or 83% savings) and a **Best Effort** reduction factor of $64 \times$ (or 98% savings).

²⁴Again, we cannot simply multiply the static saving potential with $(1/87\%)^{10years} = 4.03 \times$, for the same reason as noted higher in the text.

For both calculations, the range between the Moderate Effort and Best Effort reduction factor is striking; they are more than an order of magnitude apart. The reason is twofold. First, the reported power saving potential for a similar approach can differ substantially between different works. Consider for example some of the values in Table VIII or in Table XI. This stems from differences in the assumed baseline (topologies, architectures, traffic matrices, power models), and the way in which power savings are evaluated and reported. Often maximum saving values are reported, at other times average savings, and sometimes it is plain unclear what the baseline is for the savings. This makes both the results potentially diverse, and very hard for the reader (and the authors of the current study) to compare energy saving results across different publications. As already noted by Bianzino et al. [103], [104], there is a clear need for a consistent methodology and benchmark to evaluate different power saving techniques. As of yet, this has not yet been sufficiently addressed. However, we do think that our survey at least partly addresses his suggestion for 'the creation [...] of a green repository of energy-related figures [...] to foster future research'. This survey should make it much more easier for authors to crosscheck their results with community-wide results. The second reason for the wide range in both reduction factors is that part of our assessment is unavoidably subjective. We have tried to substantiate and justify why we chose the final reduction factors for each category as we did, and we did so to the best of our knowledge and the information we were able to find. Still, this is no 'hard science', and different authors would probably have made slightly different choices. However, we have made as much information available as we could, which should allow any reader to make his own assessment and estimation of the combined reduction factors.

B. Sensitivity to the average link length

As described in Section III-B, for our estimations we assumed that backbone networks have an average link length of 800 km. This corresponds to each link having 10 OLAs. The end result is that OLAs contribute 9% to the total power consumption in the reference backbone network (see Table II).

In Fig. 10 we show how the total static reduction factor is influenced by changing the link length to 300 km and 1300 km. As can be seen, higher average link lengths slightly decrease the power reduction factor. This is mainly driven by the contribution from applying optical bypass. When the link lengths increase, the relative weight of the OLAs to the total power consumption increases as well. As the OLAs don't contribute when applying optical bypass, the reductions of applying optical bypass are relatively decreased. For shorter link lengths, the inverse reasoning applies.

Overall, the influence of the average link length is relatively small, especially considering the level of accuracy associated with our reduction factors.

 $^{^{22}}$ The use of 2 decimal places for the reduction factors in Fig. 9(a) is to easily distinguish between the different factors, rather than being indicative for the level of accuracy of the factors.

²³Following our model in Section III-C, the combined potential can strictly speaking *not* be calculated by multiplying each of the individual factors (as reduction factors differ for different equipment types, with equipment types having different weights). However, in our case the outcome of both the correct calculation as well as the simple multiplication approximation are nearly identical. For the Moderate Effort scenario we have a reduction factor of 2.30× (correct calculation) and 2.29× (approximation), for the Best Effort scenario we have $30.9 \times$ (correct calculation) and $31.4 \times$ (approximation).





(a) Static reduction factors for each of the individual approaches (higher is better). The combined reduction potential of the power consumption in backbone networks is $2.3 \times$ (Moderate Effort) and $31 \times$ (Best Effort).

(b) Combined reduction factor potential of both the static approaches and the time-dependent 10% annual energy-per-bit reduction of telecom equipment (higher is better). By the year 2020, the combined reduction potential of the power consumption in backbone networks reaches $6.0 \times$ (Moderate Effort) and $64 \times$ (Best Effort).

Fig. 9. Reduction factors for the Moderate Effort and Best Effort scenarios for (a) the static-approaches only, and (b) the static and time-dependent approaches. Note that vertical axis uses a logarithmic scale so that the (sub)lengths are a measure for the individual reduction factors, and can easily be stacked.



Fig. 10. Sensitivity of the combined reduction factor to the average link length (with reference network hop count H = 3). Higher average link lengths slightly decrease the power reduction potential. However, the sensitivity is relatively low.



Fig. 11. Sensitivity of the combined reduction factor to the average hop count H (with reference network link length 800 km), over the range $H = 2 \rightarrow 5$. In practice, a higher hop count increases the reduction potential; this increase can be almost completely attributed to the fact that we also increased/decreased the power saving potential of employing optical bypass accordingly.

C. Sensitivity to the hop count H

The occurrence of the $(H + c_x)$ term in our model made our reduction factor calculation dependent on the average hop count H, see Section III-C. In Fig. 11 we show the influence of changing the hop count H = 3 to either 2 or 5. It is important to note that we have changed the reduction potential of optical bypass to 2 and 5 respectively (i.e., its full potential), where it was 3 initially (see Table XIV). As expected, an increase in the hop count increases the total reduction factor, and vice versa. This can almost completely be attributed to the effect of employing optical bypass; for H = 2 the reduction potential of optical bypass drops to $1.56 \times$, and for H = 5 goes up to $2.78 \times$, compared to $2.03 \times$ for H = 3.

If the reduction potential for optical bypass is not increased accordingly, e.g. as is the case for the Moderate Effort scenario (see Table XIV), the effect of the hop count on the end result is negligible, as can be seen in Fig. 11.

D. Comparison with Green Meter

To our knowledge, the 'Green Meter $v1.0'^{25}$ white paper [10] published in July 2013 by the GreenTouch consortium is the only public report that adopts a similar methodology as we do in our study. Therefore it is especially relevant to compare our findings with the Green Meter findings.

There are several important things that the report brings to the table. First, the report does not exclusively focus on the core (backbone) network, but also includes the wireline access and mobile access network. In that sense it is a more holistic and complete approach (see Table I for an overview

²⁵The Green Meter report is a work in progress, and subsequent updates and refinements to the model and estimations are planned.

 TABLE XVII

 GREENTOUCH'S GREEN METER POWER REDUCTION ESTIMATIONS [10]

Approach	Green Meter	Our study ^(a)
Improved components ^(b)	27.0 ×	-
PUE improvement $(2.0 \rightarrow 1.5)$	$1.3 \times$	$1.9 \times$
Eff. impr. (13% p.a.)	3.6 imes	$4.0 \times$
Impr. optical interconnects	$1.5 \times$	-
Improved system design and integra-	5.0 imes	-
tion		
Mixed line rates	1.2 imes	115.
Physical topology optim. to diurnal cycles	1.1 ×	} 1.5×
Optical bypass, sleep and low-energy state	1.8 ×	2.0 imes
Overall reduction factor	64.0×	-

^a We list the Best Effort reduction factors. See text for more details.

^b The net reduction potential $(27\times)$ is smaller than the $35.9\times$ obtained by simply multiplying all the individual contributing factors. This is due to the same reason as in our study, i.e., the four individual factors apply only to the IP routers and transponders, with the PUE improvement also applying to the OXCs.

of the reported power reduction factors). Second, most of the different power reduction approaches are evaluated on a consistent benchmark framework, whereas most of the papers in our survey use different power values, reference networks and optimization parameters (the work by Idzikowski et al. [81] is a notable exception). On the other hand, our focus on the backbone network extends to research outside of GreenTouch. This provides the reader with an insight in the consensus and uncertainty among the wider research community. Our report also provides both a Moderate Effort and a Best Effort scenario, which again contributes to an understanding of the range of achievable power savings. Finally, an independent approach allows to compare both works and perform a (limited) crossvalidation, as we do here.

The reduction factors reported in the Green Meter study are listed in Table XVII. The different approaches cannot exactly be mapped to our categorization, but nonetheless we can compare several approaches. The PUE improvement factor is listed as $1.3 \times$. This is nearly identical to what we assume for the Moderate Effort scenario, whereas our Best Effort scenario goes as high as $1.85 \times$. Our motivation for the latter is mainly grounded in the fact that the baseline PUE can be worse than 2.0. The $3.6 \times$ reduction factor resulting from the 13% of yearly efficiency improvement of IP routers and transponders is similar to our assumption of 15% efficiency improvement per year for the same equipment²⁶. We do not specifically account for the improved optical interconnects and the improved system design and integration; if we would, they would fit under the power rating factor $\frac{P}{C}$ (most likely improved chassis utilization and sleep modes on a short time scale) and perhaps (partly) under the overprovisioning factor. Unfortunately, the Green Meter report does not provide details or references that underpin the very large reduction potential of improved system design and integration $(5.0\times)$, which makes it hard to properly assess its potential; part of this information is likely based on proprietary vendor research and not yet

public. Likely it is an estimate of what clean-slate hardware design could bring to the table, where our focus is perhaps less aggressive and more about incremental improvements to hardware. The mixed line rate approach selects the power optimal combination of 40G, 100G and 400G interfaces for the given traffic distribution. We have not explicitly identified this as an approach, but we can consider it as similar to what we described as PAR (as PAR grooms traffic demands to more suitable-sized interfaces). However, we considered PAR in the context of exploiting daily traffic variations, which is similar to the Green Meter's physical topology optimization to diurnal cycles. Combining the savings of both approaches $(1.2 \times \text{ and }$ $1.1\times$) this gives a $1.3\times$ reduction, which is not too far off of the $1.5 \times$ reduction we estimated for sleep modes that exploit daily traffic variations. Finally, Green Meter's optical bypass approach seems to include other techniques such as sleep modes and low-power states as well. Nonetheless, this maps rather well to our estimation of $2.03 \times$ for employing optical bypass.

Approaches that are not (yet) taken into account in the Green Meter, but that we do take into account are (a) protection related measures, (b) compression, and (c) caching.

Coincidentally, the Green Meter estimate of the combined potential ($64\times$) is identical to our Best Effort reduction factor ($64\times$). This is even more surprising, as we consider some approaches that Green Meter does not, and vice versa. The approach of the Green Meter seems to be more focussed on aggressive hardware optimisation and clean-slate design, whereas our focus is slightly more towards the application layers (cf. caching). On the other hand, our Moderate Effort reduction factor ($4.9\times$) is well below the above figure.

E. The role of legacy networks

The baseline in our estimations is an IP-over-WDM network. This means we do not take into account the legacy network equipment and intermediate transport technologies such as Asynchronous Transfer Mode (ATM) and SDH that are probably²⁷ still widely present in current-day networks. As we already pointed out in Section I, this is an important limitation and point to be aware of. It implies that our power reduction estimations cannot be interpreted as applicable to current-day networks. Rather, they should be seen as an estimation for the potential to improve current state-of-the art equipment.

Nonetheless, we can try to briefly assess how our estimations can be applied to legacy equipment and networks. From the approaches we identified in Fig. 9(a), only few of them seem to be applicable to legacy equipment. Reducing the overhead power is certainly one of them, and fortunately provides good reduction potential $(1.85\times)$. Compression and caching probably can be applied to legacy networks as well, but the reduction potential $(1.23\times)$ is feeble and unsure. The other

 $^{^{26}}$ Don't confuse this with our overall backbone efficiency improvement of 13% per year, see Table XII.

²⁷We were not able to retrieve operator data on the actual distribution of deployed equipment; while part of this might be attributed to confidentiality issues, it seems not unlikely that up-to-date and network-wide inventories of such equipment are not readily available either.

approaches would very likely require hardware upgrades to support features such as sleep modes, which is unrealistic for legacy equipment. And while putting equipment to sleep on a daily time scale might in theory be feasible, it is very unlikely that operators would resort to such measures with equipment that was not designed for it.

From our very brief assessment, it appears that if the power consumption of legacy network equipment is to be reduced drastically, the best option might be to replace it altogether with new, more energy-efficient equipment.

VI. CONCLUSION AND RECOMMENDATIONS

A novel methodology to categorize and survey power saving approaches — In this paper we have surveyed a number of approaches to reduce the energy consumption of backbone IP-over-WDM telecommunication networks. The idea for this study was conceived when we noticed that in the available surveys the power saving potential (i.e. the reported saving percentages) of the works they survey was typically not discussed. This is striking, as estimating the energy saving potential of different techniques can be considered as one of the core drivers behind research in 'green networks'. The general approach in most existing surveys is a categorization of the various efficiency measures based on for example scale (e.g., optimizing at circuit-level vs. network level) or time (e.g., online sleep mode decisions vs. a preplanned energyefficient network design). In contrast, our survey approach has been based on mapping the techniques onto a concise analytical power consumption model. This has the advantage that the power reduction potential associated with each of the different techniques can be rather easily multiplied in order to be combined to a total reduction factor.

Specifically, our power saving approaches have been disaggregated in those that reduce the power rating of equipment (a measure for their efficiency), the traffic in the network (i.e., the amount of transported bits), the number of network hops, the power associated with equipment installed for traffic protection purposes, equipment installed for dealing with unexpected traffic and future traffic growth (i.e., overprovisioning), and the power associated with external overhead such as cooling and power provisioning losses.

Power reduction potential in the backbone — We assessed the power reduction potential for a Moderate Effort scenario and a Best Effort scenario, based on the uncertainty in the reported saving potential and (to a minor extent) operational and technical issues. We estimate the combined Moderate Effort reduction potential of once-only approaches at $2.3\times$, and their Best Effort reduction potential at $31\times$. If we factor in the projected yearly efficiency improvements driven by Moore's law (i.e. CMOS efficiency improvements), the 10-year reduction potential of both values are roughly (a little bit more than) doubled to $6.0\times$ and $64\times$ respectively. The large difference between the Moderate Effort and Best Effort estimate is caused mainly by the disparity and the lack of clarity of the reported savings, and our (to an extent unavoidable, subjective) assessment of the feasibility of implementation. The large difference also reflects the uncertainty we encountered when evaluating many of the surveyed works, and confirms earlier conclusions by Bianzino et al. that a standardized methodology and benchmark is deeply needed. Nonetheless we think that the general trends and findings are sound. We do provide ample information in our survey for the reader to make his own assessment.

On an individual level, a number of approaches stand out with their reduction potential. We found that the approaches that focus on bypassing power-hungry IP hops, reducing the power associated with the external overhead (i.e. improving the PUE), and putting overcapacity to sleep can reduce the power consumption most; each with roughly a factor 2 in the Best Effort scenario. Especially the latter two techniques also provide a relatively high Moderate Effort reduction potential, which is an indication for our confidence in their potential. More general, the technique of applying sleep modes shows good potential across a variety of applications.

Outrunning the traffic growth? — Considering that the traffic is projected to increase close to 10-fold in the next 10 years, the power reduction potential of our Moderate Effort scenario $(6.0\times)$ will not be sufficient to halt the power consumption increase in backbone telecommunication networks. However, the estimated $31-64\times$ reduction for our Best Effort scenario shows that enough potential is likely available to halt, and even reduce, the absolute power consumption. This will no doubt require significant efforts from an economic, technical and operational point of view. These issues have not been considered in our survey, and require extensive investigations on their own.

Recommendations for future research — From what we learned while doing this survey, we would like to make the following recommendations to our readers. For the research community: (a) Any paper on energy-efficient network techniques should clearly distinguish the various effects that result in a power consumption reduction; e.g., side effects such as an increase in the link load should be clearly identified. (b) The baseline scenario should be clearly identified; often we found that it was ambiguous whether the baseline was e.g. a common shortest-path, minimized delay scenario, or another energyoptimized scenario. (c) It should be clearly stated whether the reported saving potential applies to average or maximum savings, and to which equipment it applies (e.g., only line cards, IP equipment, or the complete network). (d) There is a need for studies that compare various power saving approaches across a consistent baseline (such as topology, power model, and traffic model). This allows for a cross-validation of reported results. For network operators, vendors and policy makers alike, we would like to state that there is really a need for a more open spirit with respect to the publication of power consumption values of network devices and complete networks. While research into energy efficiency of ICT is already several years ongoing, still very little credible power consumption values from vendors are available and a large range of equipment power consumption values are used in various works. The availability of such values directly from vendors would probably improve the consistency of the power consumption values used as an input in future research.

Finally, we hope that this survey is useful for researchers in the short-term to mid-term to check and validate their own (intermediate) power saving results across a wide set of existing works, and thereby perform a first validation of their findings.

ACKNOWLEDGMENTS

The authors would like to thank the following people for their feedback and input: Gianluca Griffa, Peter Lambert, Sofie Lambert, Christoph Lange, Esther Le Rouzic, and Andrew Lord. Of course, any errors or misinterpretations are the responsibility of the authors. The work described in this paper was carried out with the support of the iMinds-project GreenICT, and the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257740 (Network of Excellence TREND).

REFERENCES

- D. C. Kilper, G. Atkinson, S. K. Korotky, S. Goyal, P. Vetter, D. Suvakovic, and O. Blume, "Power trends in communication networks," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 17, pp. 275–284, March 2011. doi:10.1109/JSTQE.2010.2074187.
- [2] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch, "Energy consumption of telecommunication networks and related improvement options," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 17, pp. 285–295, April 2011. doi:10.1109/JSTQE.2010.2053522.
- [3] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, "Energy consumption in optical IP networks," *Journal of Lightwave Technology*, vol. 27, pp. 2391–2403, July 2009. doi:10.1109/JLT.2008.2010142.
- [4] D. Kilper, K. Guan, K. Hinton, and R. Ayre, "Energy challenges in current and future optical transmission networks," *Proceedings of the IEEE*, vol. 100, pp. 1168–1187, May 2012. doi:10.1109/JPROC.2012.2186102.
- [5] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, "Energy efficiency in telecom optical networks," *Communications Surveys & Tutorials, IEEE*, vol. 12, pp. 441–458, November 2010. doi:10.1109/SURV.2011.062410.00034.
- [6] R. Bolla, R. Bruschi, F. Davoli, and F. Cucchietti, "Energy efficiency in the future internet: a survey of existing approaches and trends in energy-aware fixed network infrastructures," *Communications Surveys & Tutorials, IEEE*, vol. 13, pp. 223–244, May 2011. doi:10.1109/SURV.2011.071410.00073.
- [7] A. P. Bianzino, C. Chaudet, D. Rossi, and J.-L. Rougier, "A survey of green networking research," *Communications Surveys & Tutorials, IEEE*, vol. 14, pp. 3–20, February 2012. doi:10.1109/SURV.2011.113010.00106.
- [8] S. Roy, "Energy logic for telecommunications." Emerson, White Paper, 2008.
- [9] M. Parker and S. Walker, "Roadmapping ICT: an absolute energy efficiency metric," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 3, pp. A49–A58, July 2011. doi:10.1364/JOCN.3.000A49.
- [10] GreenTouch, "GreenTouch Green Meter research study: Reducing the net energy consumption in communications networks by up to 90% by 2020." http://www.greentouch.org/uploads/documents/GreenTouch_ Green_Meter_Research_Study_26_June_2013.pdf, June 2013.
- [11] E. Le Rouzic, E. Bonetto, L. Chiaraviglio, F. Giroire, F. Idzikowski, F. Jimenez, C. Lange, J. Montalvo, F. Musumeci, I. Tahiri, A. Valenti, W. Van Heddeghem, Y. Ye, A. Bianco, and A. Pattavina, "TREND towards more energy-efficient optical networks," in *ONDM*, (Brest, France), pp. 211–216, April 2013.
- [12] E. Le Rouzic, R.-M. Indre, L. Chiaraviglio, F. Musumeci, A. Pattavina, J. Lopez Vizcaino, Y. Ye, W. Van Heddeghem, A. Bianco, E. Bonetto, M. Meo, F. Jimenez, F. Idzikowski, and R. Cuevas, "TREND big picture on energy-efficient backbone networks," in *Digital Communications-Green ICT (TIWDC), 2013 24th Tyrrhenian International Workshop on*, pp. 1–6, IEEE, September 2013. doi:10.1109/TIWDC.2013.6664209.

- [13] W. Van Heddeghem, M. C. Parker, S. Lambert, W. Vereecken, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester, "Using an analytical power model to survey power saving approaches in backbone networks," in *Networks and Optical Communications (NOC), 2012 17th European Conference on*, (Vilanova i la Geltru, Spain), pp. 1–6, IEEE, June 2012. doi:10.1109/NOC.2012.6249942.
- [14] J. Malmodin, A. Moberg, D. Lundén, G. Finnveden, and N. Lövehagen, "Greenhouse gas emissions and operational electricity use in the ICT, and entertainment & media sectors," *J. Ind. Ecol.*, vol. 14, pp. 770–790, October 2010. doi:10.1111/j.1530-9290.2010.00278.x.
- [15] S. Lambert, W. Van Heddeghem, W. Vereecken, B. Lannoo, D. Colle, and M. Pickavet, "Worldwide electricity consumption of communication networks," *Optics Express*, vol. 20, pp. B513–B524, December 2012. doi:10.1364/OE.20.00B513.
- [16] W. Van Heddeghem, F. Idzikowski, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, "Power consumption modeling in optical multilayer networks," *Photonic Network Communications*, January 2012. doi:10.1007/s11107-011-0370-7.
- [17] J. C. C. Restrepo, C. G. Gruber, and C. M. Machuca, "Energy profile aware routing," in *Communications Workshops*, 2009. *ICC Workshops* 2009. *IEEE International Conference on*, (Dresden, Germany), pp. 1–5, IEEE, June 2009. doi:10.1109/ICCW.2009.5208041.
- [18] R. Bolla, R. Bruschi, C. Lombardo, and D. Suino, "Evaluating the energy-awareness of future internet devices," in *High Performance Switching and Routing (HPSR), 2011 IEEE 12th International Conference on*, pp. 36–43, July 2011. doi:10.1109/HPSR.2011.5986001.
- [19] E. R. Masanet, R. E. Brown, A. Shehabi, J. G. Koomey, and B. Nordman, "Estimating the energy use and efficiency potential of US data centers," *Proceedings of the IEEE*, vol. 99, pp. 1440–1453, August 2011. doi:10.1109/JPROC.2011.2155610.
- [20] W. Van Heddeghem, F. Idzikowski, E. Le Rouzic, J. Mazeas, H. Poignant, S. Salaun, B. Lannoo, and D. Colle, "Evaluation of power rating of core network equipment in practical deployments," in *Online Conference on Green Communications (GreenCom)*, 2012 IEEE, pp. 126–132, September 2012. doi:10.1109/GreenCom.2012.6519628.
- [21] R. Hülsermann, A. Betker, M. Jäger, S. Bodamer, M. Barry, J. Späth, C. Gauger, and M. Köhn, "A set of typical transport network scenarios for network modelling," *ITG FACHBERICHT*, no. 182, pp. 65–72, 2004.
- [22] C. Belady, A. Rawson, J. Pfleuger, and T. Cader, "Green Grid data center power efficiency metrics: PUE and DCIE," White Paper #6, The Green Grid, 2008.
- [23] V. Avelar, D. Azevedo, and A. French, "PUE: A comprehensive examination of the metric," White Paper #49, The Green Grid, 2012.
- [24] EPA, "Report to congress on server and data center energy efficiency." U.S. Environmental Protection Agency, http://www.energystar.gov/ia/partners/prod_development/downloads/ EPA_Datacenter_Report_Congress_Final1.pdf, 2007.
- [25] Google, "Data center efficiency power usage effectiveness." http://www.google.com/about/datacenters/inside/efficiency/ power-usage.html, 2012.
- [26] M. Stansberry and J. Kudritzki, "Uptime Institute 2012 data center industry survey." Uptime Institute, 2012. Available at: http: //uptimeinstitute.com/2012-survey-results.
- [27] 42u, "What is PUE? what is DCiE?." http://www.42u.com/measurement/ pue-dcie.htm#What-is-PUE?
- [28] N. Rasmussen, Implementing Energy Efficient Data Centers (White Paper 114, Revision 1). APC, 2011.
- [29] Intel, "The state of data center cooling a review of current air and liquid cooling solutions." ftp://download.intel.com/technology/EEP/ data-center-efficiency/state-of-date-center-cooling.pdf, 2008.
- [30] Google, "Data center best practices." http://www.google.com/about/ datacenters/best-practices.html, 2012.
- [31] 42u, "Cooling solutions for server racks, computer rooms, server rooms and data centers." http://www.42u.com/42u-rack-cooling.htm.
- [32] E. Pinheiro, W. Weber, and L. Barroso, "Failure trends in a large disk drive population," in *Proceedings of the 5th USENIX conference on File and Storage Technologies*, pp. 2–2, 2007.
- [33] S. Greenberg, E. Mills, B. Tschudi, P. Rumsey, and B. Myatt, "Best practices for data centers: Lessons learned from benchmarking 22 data centers," *Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings in Asilomar, CA. ACEEE*, vol. 3, pp. 76–87, August 2006.
- [34] X. Fan, W. Weber, and L. Barroso, "Power provisioning for a warehousesized computer," ACM SIGARCH Computer Architecture News, vol. 35, pp. 13–23, May 2007. doi:10.1145/1273440.1250665.
- [35] J. Koomey, "The NYT article on power, pollution, and the internet: My initial comments," September 2012. Blog post available at: http: //www.koomey.com/post/32281701993.

- [36] W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester, "Trends in worldwide ICT electricity consumption from 2007 to 2012," *Computer Communications*, vol. 50, pp. 64–76, September 2014. doi:10.1016/j.comcom.2014.02.008.
- [37] "Responsabilité sociale d'entreprise, rapport détaillé." France Telecom - Orange, 2011. Available at: http://www.orange.com/fr/finance/nbsp4/ assemblee-generale-2013/rapports-annuels.
- [38] P. Blasco, "Telefónica inaugura en Alcalá de Henares uno de los mayores centros de datos [Telefonica opens in Alcala de Henares one of the largest data centers]." El Mundo.es, April 30, 2013. Available at: http://www.elmundo.es/elmundo/2013/04/30/navegante/1367304168. html.
- [39] "The 2010/2011 corporate responsibility report." Deutsche Telekom. Retrieved on 2013-11-28 at: http://www.cr-report.telekom.com/ site11/en/klima-und-umwelt/umweltfreundliche-netzinfrastruktur/ klimatechnik-und-energieeinsatz/index.php?page=1#ff,43918,43917, 43916,43915,43914,43913,43920,43919.
- [40] "Corporate responsibility report 2012." Deutsche Telekom. Retrieved on 2013-11-28 at: http://www.cr-report.telekom.com/site13/sites/default/ files/pdf/2012_corporate_responsibility_report-individuell.pdf.
- [41] S. Le Masson and D. Nortershauser, "Complex wall for indirect freecooling in datacentres," in *Telecommunications Energy Conference* (*INTELEC*), 2012 IEEE 34th International, (Scottsdale, USA), pp. 1–6, IEEE, October 2012. doi:10.1109/INTLEC.2012.6374506.
- [42] C. Cavdar, F. Buzluca, and L. Wosinska, "Energy-efficient design of survivable WDM networks with shared backup," in *Global Telecommu*nications Conference (GLOBECOM 2010), 2010 IEEE, (Miami, USA), pp. 1–5, IEEE, December 2010. doi:10.1109/GLOCOM.2010.5683203.
- [43] A. Jirattigalachote, C. Cavdar, P. Monti, L. Wosinska, and A. Tzanakaki, "Dynamic provisioning strategies for energy efficient WDM networks with dedicated path protection," *Optical Switching and Networking*, vol. 8, pp. 201–213, July 2011. doi:10.1016/j.osn.2011.03.008.
- [44] A. Coiro, M. Listanti, A. Valenti, and F. Matera, "Reducing power consumption in wavelength routed networks by selective switch off of optical links," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 17, pp. 428–436, April 2011. doi:10.1109/JSTQE.2010.2052791.
- [45] N.-H. Bao, L.-M. Li, H.-F. Yu, Z.-Z. Zhang, and H.-B. Luo, "Poweraware provisioning strategy with shared path protection in optical WDM networks," *Optical Fiber Technology*, vol. 18, pp. 81–87, March 2012. doi:10.1016/j.yofte.2011.12.003.
- [46] F. Musumeci, M. Tornatore, J. Lopez Vizcaino, Y. Ye, and A. Pattavina, "Energy-efficiency of protected IP-over-WDM networks with sleepmode devices," *Journal of High Speed Networks*, vol. 19, pp. 19–32, March 2013. doi:10.3233/JHS-130463.
- [47] J. Lopez Vizcaino, Y. Ye, V. Lopez, F. Jimenez, R. Duque, F. Musumeci, A. Pattavina, and P. Krummrich, "Differentiated quality of protection to improve energy efficiency of survivable optical transport networks," in *Optical Fiber Communication Conference*, (Anaheim, USA), Optical Society of America, March 2013. doi:10.1364/OFC.2013.OM3A.1.
- [48] O. Tamm, C. Hermsmeyer, and A. M. Rush, "Eco-sustainable system and network architectures for future transport networks," *Bell Labs Technical Journal*, vol. 14, pp. 311–327, Winter 2010. doi:10.1002/bltj.20418.
- [49] X. Dong, A. Lawey, T. E. El-Gorashi, and J. M. Elmirghani, "Energyefficient core networks," in *Optical Network Design and Modeling* (*ONDM*), 2012 16th International Conference on, (Colchester, UK), pp. 1–9, IEEE, April 2012. doi:10.1109/ONDM.2012.6210196.
- [50] R. Kothiyal, V. Tarasov, P. Sehgal, and E. Zadok, "Energy and performance evaluation of lossless file data compression on server systems," in *Proceedings of SYSTOR 2009: The Israeli Experimental Systems Conference*, (Haifa, Israel), p. 4, ACM, May 2009. doi:10.1145/1534530.1534536.
- [51] U. Lee, I. Rimac, and V. Hilt, "Greening the internet with content-centric networking," in *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking (e-Energy '10)*, (Passau, Germany), pp. 179–182, ACM, April 2010. doi:10.1145/1791314.1791342.
- [52] U. Lee, I. Rimac, D. Kilper, and V. Hilt, "Toward energy-efficient content dissemination," *Network, IEEE*, vol. 25, pp. 14–19, April 2011. doi:10.1109/MNET.2011.5730523.
- [53] N. Osman, T. El-Gorashi, and J. Elmirghani, "Reduction of energy consumption of video-on-demand services using cache size optimization," in *Wireless and Optical Communications Networks (WOCN), 2011 Eighth International Conference on*, (Paris, France), pp. 1–5, IEEE, May 2011. doi:10.1109/WOCN.2011.5872923.
- [54] N. I. Osman, T. El-Gorashi, and J. M. Elmirghani, "Caching in green

IP over WDM networks," *Journal of High Speed Networks*, vol. 19, pp. 33–53, March 2013. doi:10.3233/JHS-120458.

- [55] L. Chiaraviglio and I. Matta, "An energy-aware distributed approach for content and network management," in *Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on*, (Shanghai, China), pp. 337–342, IEEE, April 2011. doi:10.1109/INFCOMW.2011.5928834.
- [56] L. Chiaraviglio and I. Matta, "Energy-aware network management and content distribution," in *Green Communications: Theoretical Fundamentals, Algorithms, and Applications* (J. Wu, S. Rangan, and H. Zhang, eds.), CRC Press, 2012.
- [57] R. Modrzejewski, L. Chiaraviglio, I. Tahiri, F. Giroire, E. Le Rouzic, E. Bonetto, F. Musumeci, R. Gonzalez, and C. Guerrero, "Energy efficient content distribution in an ISP network," in *Global Communications Conference (GLOBECOM), 2013 IEEE*, (Atlanta, USA), December 2013. doi:10.1109/GLOCOM.2013.6831508.
- [58] U. Mandal, M. F. Habib, S. Zhang, C. Lange, A. Gladisch, and B. Mukherjee, "Adopting hybrid CDN-P2P in IP-over-WDM networks: an energy-efficiency perspective," *Journal of Optical Communications and Networking*, vol. 6, pp. 303–314, February 2014. doi:10.1364/JOCN.6.000303.
- [59] Cisco, "Cisco visual networking index: Forecast and methodology, 2010–2015," 2011.
- [60] Cisco, "Global IP traffic forecast and methodology, 2006-2011," 2007.
- [61] Cisco, "Cisco visual networking index: Forecast and methodology, 2008–2013," 2009.
- [62] Cisco, "Cisco visual networking index: Forecast and methodology, 2012–2017," 2013.
- [63] Y. Wu, P. A. Chou, and S.-Y. Kung, "Minimum-energy multicast in mobile ad hoc networks using network coding," *Communications, IEEE Transactions on*, vol. 53, pp. 1906–1918, November 2005. doi:10.1109/TCOMM.2005.857148.
- [64] X. Liu, K. Fouli, R. Kang, and M. Maier, "Network-coding-based energy management for next-generation passive optical networks," *Journal of Lightwave Technology*, vol. 30, pp. 864–875, March 2012. doi:10.1109/JLT.2011.2182633.
- [65] GreenTouch, "GreenTouch 2012–2013 annual report." http://www.greentouch.org/uploads/documents/GreenTouch_2013_ Annual_Report.pdf, June 2013.
- [66] J. Koomey, S. Berard, M. Sanchez, and H. Wong, "Implications of historical trends in the electrical efficiency of computing," *Annals of the History of Computing, IEEE*, vol. 33, pp. 46–54, March 2011. doi:10.1109/MAHC.2010.28.
- [67] T. Ebrahimi, "What future for image/video compression." http://www. icme2011.org/files/icme2011_panel1_ebrahimi.pdf, July 2011.
- [68] G. Heydon, "IPTV (is changing fast)." ITU-AIBD-ABU Regional Workshop – Strengthening Digital Broadcasting Experience, Hanoi (Vietnam), available at: http://www.itu.int/ITU-D/asp/CMS/Events/2011/ DigitalBroadcast-May2011/S8_GeofHeydon.pdf, May 2011.
- [69] G. Peng, "CDN: Content Distribution Network," Arxiv preprint cs/0411069, 2004.
- [70] A. Feldmann, A. Gladisch, M. Kind, C. Lange, G. Smaragdakis, and F. Westphal, "Energy trade-offs among content delivery architectures," in *Telecommunications Internet and Media Techno Economics (CTTE)*, 2010 9th Conference on, (Ghent, Belgium), pp. 1–6, IEEE, June 2010. doi:10.1109/CTTE.2010.5557700.
- [71] D. T. Neilson, "Photonics for switching and routing," *Selected topics in quantum electronics, IEEE Journal of*, vol. 12, pp. 669–678, August 2006. doi:10.1109/JSTQE.2006.876315.
- [72] S. Han, "Moore's law and energy and operations savings in the evolution of optical transport platforms," *Communications Magazine, IEEE*, vol. 48, pp. 66–69, February 2010. doi:10.1109/MCOM.2010.5402665.
- [73] L. Ceuppens, "Planning for energy efficiency: Networking in numbers," in OFC/NFOEC 2009, (San Diego, USA), 2009.
- [74] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," in *Proceedings of INFOCOM*, (Phoenix, USA), pp. 1130–1138, April 2008. doi:10.1109/INFOCOM.2008.93.
- [75] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "A power benchmarking framework for network devices," in *Proceedings of NETWORKING, Aachen, Heidelberg*, (Aachen, Germany), pp. 795–808, 2009. doi:10.1007/978-3-642-01399-7_62.
- [76] R. Tucker, "Green optical communications part II: Energy limitations in networks," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 17, pp. 261–274, April 2011. doi:10.1109/JSTQE.2010.2051217.

- [77] M. C. Parker, S. D. Walker, R. Martin, W. Van Heddeghem, and B. Lannoo, "Energy-efficient master-slave edge-router upgrade paths in active remote nodes of next-generation optical access," in *Asia Communications and Photonics Conference*, (Guangzhou, China), Optical Society of America, November 2012.
- [78] C. Lange and A. Gladisch, "Energy efficiency limits of load adaptive networks," in *Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference, 2010 Conference on* (OFC/NFOEC), (San Diego, USA), pp. 1–3, IEEE, March 2010. doi:10.1364/OFC.2010.OWY2.
- [79] L. Liu and B. Ramamurthy, "Rightsizing bundle link capacities for energy savings in the core network," in *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE*, (Houston, USA), pp. 1– 6, IEEE, December 2011. doi:10.1109/GLOCOM.2011.6134490.
- [80] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz, "Dynamic routing at different layers in IP-over-WDM networks maximizing energy savings," *Optical Switching and Networking*, vol. 8, pp. 181–200, July 2011. doi:10.1016/j.osn.2011.03.007.
- [81] F. Idzikowski, E. Bonetto, L. Chiaraviglio, A. Cianfrani, A. Coiro, R. Duque, F. Jiménez, E. Le Rouzic, F. Musumeci, W. Van Heddeghem, J. López Vizcaíno, and Y. Ye, "TREND in energy-aware adaptive routing solutions," *Communications Magazine, IEEE*, November 2013. doi:10.1109/MCOM.2013.6658659.
- [82] Y. Zhang, M. Tornatore, P. Chowdhury, and B. Mukherjee, "Timeaware energy conservation in IP-over-WDM networks," in *Photonics in Switching*, Optical Society of America, July 2010. doi:10.1364/PS.2010.PTuB2.
- [83] W. Fisher, M. Suchara, and J. Rexford, "Greening backbone networks: Reducing energy consumption by shutting off cables in bundled links," in *Proceedings of the first ACM SIGCOMM workshop on Green networking*, pp. 29–34, ACM, 2010. doi:10.1145/1851290.1851297.
- [84] F. Idzikowski, L. Chiaraviglio, R. Duque, F. Jiménez, and E. Le Rouzic, "Green horizon: Looking at backbone networks in 2020 from the perspective of network operators," in *Communications (ICC), 2013 IEEE International Conference on*, (Budapest, Hungary), pp. 4455– 4460, IEEE, June 2013. doi:10.1109/ICC.2013.6655268.
- [85] L. Chiaraviglio, M. Mellia, and F. Neri, "Minimizing ISP network energy cost: Formulation and solutions," *IEEE/ACM Transactions on Networking (TON)*, vol. 20, pp. 463–476, April 2012. doi:10.1109/TNET.2011.2161487.
- [86] E. Bonetto, L. Chiaraviglio, F. Idzikowski, and E. Le Rouzic, "Algorithms for the multi-period power-aware logical topology design with reconfiguration costs," *Journal of Optical Communications and Networking*, vol. 5, pp. 394–410, April 2013. doi:10.1364/JOCN.5.000394.
- [87] R. Bolla, R. Bruschi, A. Cianfrani, and M. Listanti, "Introducing standby capabilities into next-generation network devices," in *Proceed*ings of the PRESTO, Philadelphia, USA, November 2010.
- [88] A. Coiro, M. Listanti, A. Valenti, and F. Matera, "Energy-aware traffic engineering: a routing-based distributed solution for connection-oriented IP networks," *Computer Networks*, vol. 57, pp. 2004–2020, June 2013. doi:10.1016/j.comnet.2013.03.017.
- [89] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing network energy consumption via sleeping and rateadaptation," in *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, 2008.
- [90] P. Reviriego, D. Larrabeiti, and J. A. Maestro, "Would energy efficient ethernet be effective on 10gbps optical links?," in *Photonics in Switching*, Optical Society of America, July 2010. doi:10.1364/PS.2010.PTuB5.
- [91] P. Reviriego, B. Huiszoon, V. López, R. Coenen, J. Hernández, and J. Maestro, "Improving energy efficiency in IEEE 802.3 ba high-rate ethernet optical links," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 17, pp. 419–427, April 2011. doi:10.1109/JSTQE.2010.2050136.
- [92] W. Vereecken, W. Van Heddeghem, M. Deruyck, B. Puype, B. Lannoo, W. Joseph, D. Colle, L. Martens, and P. Demeester, "Power consumption in telecommunication networks: Overview and reduction strategies," *Communications Magazine, IEEE*, vol. 49, pp. 62–69, June 2011. doi:10.1109/MCOM.2011.5783986.
- [93] F. Ilchmann *et al.*, "Efficient and optimized network architecture: Requirements and reference scenarios," Deliverable D2.1, STRONGEST Project, September 2010.
- [94] Labovitz, Craig, "What Europeans do at night." ARBOR Networks, blog post at http://www.arbornetworks.com/asert/2009/08/ what-europeans-do-at-night/, August 2009.
- [95] R. Doverspike, K. K. Ramakrishnan, and C. Chase, "Structural overview of ISP networks," in *Guide to Reliable Internet Services and Applica-*

tions (C. Kalmanek, S. Misra, and R. Yang, eds.), pp. 19–93, Springer, 2010. doi:10.1007/978-1-84882-828-5_2.

- [96] W. Van Heddeghem, F. Musumeci, F. Idzikowski, A. Pattavina, B. Lannoo, D. Colle, and M. Pickavet, "Power consumption evaluation of circuit-switched versus packet-switched optical backbone networks," in *Online Conference on Green Communications (GreenCom)*, 2013 IEEE, pp. 56–63, October 2013. doi:10.1109/OnlineGreenCom.2013.6731029.
- [97] R. Zhang-Shen and N. McKeown, "Designing a predictable internet backbone network," in *Proceedings of HotNets III*, (San Diego, USA), 2004.
- [98] E. Tego, F. Idzikowski, L. Chiaraviglio, A. Coiro, and F. Matera, "Facing the reality: Validation of energy saving mechanisms on a testbed," *Journal of Electrical and Computer Engineering*, vol. 2014, 2014. doi:10.1155/2014/806960.
- [99] M. Gupta and S. Singh, "Using low-power modes for energy conservation in ethernet lans," in *INFOCOM*, vol. 7, pp. 2451–55, 2007.
- [100] P. Reviriego, J. Hernández, D. Larrabeiti, and J. A. Maestro, "Performance evaluation of energy efficient ethernet," *Communications Letters, IEEE*, vol. 13, pp. 697–699, September 2009. doi:10.1109/LCOMM.2009.090880.
- [101] G. Shen and R. Tucker, "Energy-minimized design for IP over WDM networks," *Journal of Optical Communications and Networking*, vol. 1, pp. 176–186, June 2009. doi:10.1364/JOCN.1.000176.
- [102] Y. Huang and R. Guerin, "Does over-provisioning become more or less efficient as networks grow larger?," in *Network Protocols*, 2005. *ICNP 2005. 13th IEEE International Conference on*, (Boston, USA), pp. 11–pp, IEEE, November 2005. doi:10.1109/ICNP.2005.12.
- [103] A. P. Bianzino, A. K. Raju, and D. Rossi, "Apples-to-apples: a framework analysis for energy-efficiency in networks," *SIGMET-RICS Perform. Eval. Rev.*, vol. 38, pp. 81–85, January 2011. doi:10.1145/1925019.1925036.
- [104] A. P. Bianzino, A. Raju, and D. Rossi, "Power profiling the internet core: A sensitivity analysis," *Energy-aware Systems and Networking for Sustainable Initiatives*, p. 287, 2012. doi:10.4018/978-1-4666-1842-8.ch013.