

Using an Analytical Power Model to Survey Power Saving Approaches in Backbone Networks

(Invited Paper)

Ward Van Heddeghem*, Michael C. Parker†, Sofie Lambert*, Willem Vereecken*,
Bart Lannoo*, Didier Colle*, Mario Pickavet* and Piet Demeester*

*Department of Information Technology (INTEC) of Ghent University - IBBT, Gaston Crommenlaan 8, 9050 Gent, Belgium

†School of Computer Science and Electronic Engineering, University of Essex, Wivenhoe Park, Colchester, CO3 4HG, UK

Email: *firstname.lastname@intec.ugent.be, †mcpark@essex.ac.uk

Abstract—Several approaches exist to categorize power savings in Information and Communication Technologies (ICT) networks. In this paper we survey current power saving techniques in backbone telecommunication networks, based on an analytical power model. This categorization allows for a more intuitive understanding of the power saving potential of different approaches, and shows that in order to achieve large savings (i.e. more than 10 times reduction of the current power consumption), effort will need to be concentrated on those techniques that reduce either the traffic volume or the equipment power rating. The survey in this paper also provides the basis for a more quantitative evaluation of future power saving techniques.

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 32% currently, this reduced growth still outperforms the annual 10% efficiency increase of telecommunication equipment in the backbone [1], [2]. As such, the power consumed by Information and Communication Technologies (ICT) 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. Similarly, this is reflected in the increasing number of publications and research on this topic from academia, industry and governmental bodies alike. The research 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.

Global network power consumption estimation is best done via top-down approach, whereas evaluating power saving potentials is conversely done using a bottom-up approach — For research focusing on the latter topic (i.e., novel power saving techniques), power consumption is often 'bottom-up' estimated by getting a count of all the considered equipment, and then multiplying them with their respective power consumption values. This is in contrast with attempts to estimate

global network power consumption. Here, estimating the power consumption of an existing, large network – like the Internet – is rather more complex than might appear at first sight, and is due to a number of reasons. First, such large networks are typically not just one network, but rather a collection of networks, with different topologies, and owned and controlled by different players. Second, there are a variety of architectures deployed, often historically added-to and strung together to provide compatibility with legacy equipment. For large-scale networks it is not easy, even for operators, to estimate with reasonable detail all of the equipment involved. Third, even if a reasonable estimate can be made of equipment type and count, getting representative equipment power consumption values under representative operating conditions is not straightforward either. A major reason for this is that many equipment vendors still only disclose little information on power consumption; a few notable exceptions not withstanding. As a result, research on global power consumption estimation often uses a top-down approach, e.g. [3], instead of the bottom-up approach outlined above for evaluating novel power saving techniques.

This paper uses a top-down power model to survey opportunities for reducing power consumption — Our earlier work [4] was a response to the issue outlined above where representative power consumption values are hard to obtain. We presented best-effort power consumption reference values for the main building blocks of backbone network equipment. In addition, we also provided an analytical model to estimate, using a top-down approach, the power consumption across the various layers, such as IP switches, optical switches and transponders. This model is in-line with two earlier proposed models [3], [1], with the latter being referred to as a transactional model. In this paper, we want to use the general form of our analytical model to survey a number of power saving approaches for backbone networks. In this way, we are trying to bring together techniques for both global network power consumption estimation and evaluation of power saving, which, as we have argued, have previously been approached differently, i.e. using top-down and bottom-up approaches respectively. Using a simple analytical model also offers an intuitive approach, as the impact of certain power reducing techniques can then be (roughly) estimated given a number

of reference values for the model's parameters. The work in this paper can therefore be used as a basis for quantifying the numerical power saving potential of each technique.

Related work — For some of the power reducing techniques considered in this paper, we base ourselves on earlier work from one of the authors, i.e. [5]. Next to proposing an absolute energy efficiency metric ($\text{dB}\epsilon$) for any ICT system, the cited work 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. There are numerous other works that list and categorize power saving approaches; we only name a few below. In [6], 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 [7], and the earlier survey paper [8], the approaches are classified 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). All these categorizations come with their own merits and drawbacks. As we have said before, we feel that using an actual power model for discussing power saving techniques allows for a more insightful estimation, and can in a second step be used to quantitatively assess the potential power savings.

For the remainder of this paper, we first briefly summarize the analytical power model from [4] in Section II, before using this model in Section III to discuss a number of power reduction approaches.

II. REFERENCE POWER VALUES AND AN ANALYTICAL POWER MODEL

In this section, we briefly present the reference power values for backbone network equipment, and the analytical power model from [4].

A. Reference power values

One of the main goals of our earlier work [4] was to provide representative and consistent power consumption values for backbone equipment, by a homogenization process. The final representative power consumption value depends on the typical (rather than maximum) power consumption, as well as the appropriate power consumption values associated with the chassis and control overhead powering. In addition, the representative power consumption value is that as associated with a 100% filled rack/shelf. Finally, the representative power consumption value is also as appropriate for bidirectional equipment (i.e. full-duplex communication).

TABLE I
A ILLUSTRATIVE SELECTION OF POWER RATING VALUES FROM [4]

| Type | Power | Power rating |
|--|-------|--------------|
| Backbone IP router (optimal filling) | - | 6 W/Gbps |
| Backbone IP router (real-life filling) | - | 10 W/Gbps |
| Transponder 10G | 50 W | 5 W/Gbps |
| OLA 80 km (optimal filling) | 110 W | - |
| OLA 80 km (real-life filling) | 165 W | - |

For illustrative purposes, we list a number of values in Table I. If applicable, the table also provides power rating values, expressing the power per capacity (in W/Gbps). For a complete listing – including also Ethernet, Optical Transport Networking (OTN), and Wavelength Division Multiplexing (WDM) equipment like Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Optical Cross-Connects (OXC) etc. – please refer to [4]. A detailed technical report with source data is available as well [9].

We should also point out that, in a separate work [10], we have already explored the impact of real-life equipment filling levels on the optimal power ratings given previously. Our initial findings suggest that, as a result of sub-optimal filling, actual power ratings might be almost two times higher (i.e. worse) for IP routers, and 1.5 times higher for Optical Line Amplifiers (OLAs). We have also listed these adjusted values in Table I. We will come back to the influence of equipment filling in Section III-D.

B. Analytical power model

Our analytical model gives the total power $P_{backbone}$ in an IP-over-WDM network as the sum of the power consumption in the constituting layers:

$$\begin{aligned} P_{backbone} &= P_{ip} + P_{wdm} & (1) \\ &= P_{ip} + (P_{optsw} + P_{trs} + P_{amps} + P_{regens}), & (2) \end{aligned}$$

with P_{optsw} , P_{trs} , P_{amps} and P_{regens} respectively being the power consumption of the optical switching equipment (i.e., WDM terminals, ROADMs or OXC) the WDM transponders, the OLA power consumption and the 3R regeneration. In this paper we leave out other potential intermediate switching technologies, such as OTN and Ethernet. For more information, please refer to [4].

Conveniently, the power consumption of each layer or component is further given according to the following general form, ignoring a few minor layer-specific terms and factors:

$$P_x = \eta_{eo} \cdot \eta_{pr} \cdot N_d \cdot \overline{DC} \cdot \left(\frac{P_X}{C_X} \cdot H \right). \quad (3)$$

The external overhead factor η_{eo} accounts for the power consumption due to external cooling and facility overheads in telecom centers, with typically $\eta_{eo} = 2$. The protection factor η_{pr} accounts for traffic protection, with $\eta_{pr} = 2$ for backbone networks using a 1+1 protection scheme (i.e., all

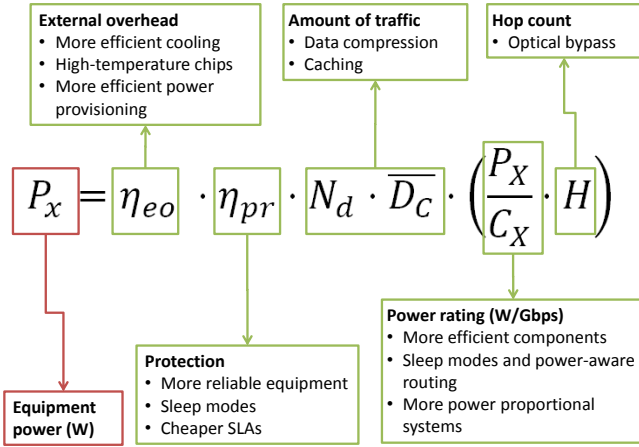


Fig. 1. Overview of power reduction approaches mapped to the general form of our analytical model

traffic is routed twice on link-disjoint paths). The combined factor $N_d \cdot \overline{D_C}$ gives the total amount of traffic in the network as a product of the number of demands N_d and the average required capacity per demand $\overline{D_C}$. The power rating factor $\frac{P_x}{C_x}$ expresses the average power per capacity (in W/Gbps) for a given equipment X , such as an IP router or transponder. Finally, the hop count H is the average number of hops between processing elements in the respective layer.

III. 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 equation (3), with a graphical overview given in Fig. 1.

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 consumed, and typically has a value of 2 [11]. 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 lower PUE values are possible¹, but this is not yet commonplace. On average, this overhead is made up of two main contributing components [13][14]: cooling and air conditioning, and efficiency losses in power provisioning.

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

¹For example, Google is claiming to have reached an annualised average for all their data centers of 1.14 by the end of 2011 [12].

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 overviews are available in [15], [16] and [17]. 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 [16]. For example, it has been observed that contrary to popular belief, hard disks do not become less reliable when running at higher temperatures [18]. 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 [5]. 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 1/5 to 1/3 of the external overhead power [14], mainly through inefficiencies in Uninterruptible Power Supply (UPS) units and Power Distribution Units (PDUs) [13]. 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 [19]. 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) of the ICT equipment, which can be substantially higher than the actual peak power [20]. Right-sized provisioning of equipment to the actual peak power of the ICT equipment would reduce the overhead power consumption.

B. Protection factor η_{pr}

The protection factor η_{pr} accounts for traffic protection. Traffic protection is typically employed in backbone networks to achieve high reliability to meet costly Service Level Agreements (SLAs). A common protection scheme in backbone networks is 1+1 protection, whereby for each demand between a source and destination node two link-disjoint IP connections are set up. This results in a protection factor $\eta_{pr} = 2$. Other protection schemes exist, such as 1:1, 1:2 and N:M protection schemes [21].

We consider the following approaches to bring down the protection factor: more reliable equipment, sleep modes, and cheaper SLAs.

More reliable equipment — In [5] it is argued that energy savings of up to 50% of the total network power consumption could be achieved by avoiding the above described dualling. This would be made possible through increased reliability

of network devices, systems and subsystems, 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.

Sleep modes — 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. Using this approach, [22] reports a potential 30% reduction of total network energy consumption. For realistic application, this would require fast sleep and wakeup times. We will return to sleep modes in Section III-D.

Cheaper SLAs — 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.

C. Amount of traffic

The amount of traffic in the network is given by the product $N_d \cdot \overline{D_C}$, i.e. the average demand $\overline{D_C}$ times the number of demands N_d .

What are potential approaches to bring down the amount of traffic in the backbone network? We will look at data compression and caching.

Data compression — Data compression, or source coding², encodes information in such a way that it requires fewer bits than in the original representation. The reference [5] estimates 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 HDTV broadcasts). Second, compression and decompression at transmit and receiving side comes at a processing, and thus energy, cost. That said, the influence of this could be minor and drop with continuously more efficient Digital Signal Processings (DSPs) units [5], [23]. Third, any advances in compression techniques will probably be canceled out by encoding more information and new modalities (such as stereoscopic view, or a higher dynamic range for audio or video) into bitstreams (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 [24]; the newest upcoming standard (HEVC), is again expected to continue this trend.

Caching — Another technique to decrease the amount of

²Source coding should not be confused with *channel coding* which in effect adds redundancy (instead of removing redundancy, as source coding does), to improve the bit error rate on noisy communication channels.

traffic on 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 [25]. 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, in [26] it is estimated that with optimal cache sizes, a reduction of up to about 40% of the total power consumption (network and caching) can be achieved.

D. Power rating

The power rating factor P/C expresses equipment power consumption as the power per unit capacity, e.g. 5 W/Gbps.

How can the power ratings of various items of equipment, such as IP routers and transponders, be reduced? We consider inherent energy efficiency improvements, sleep modes, and the more general form of power proportionality.

More efficient components — Telecommunication equipment becomes more power efficient each year. However, when discussing improvements in the power efficiency of network equipment, it is important to distinguish between electronic equipment and optical equipment. Electronic backbone network equipment has historically been improving by around 20% per year; however slowing down to about 10% per year more recently [1]. In the same work, it has been argued that this efficiency improvement might further slow down as a result of practical limitations inherent to CMOS transistor design. In [23], on the other hand, it is argued 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. Optical equipment on the other hand, does not benefit from the above efficiency trend, and tends to improve at much slower rates. This is taken into account by assuming an optical fraction OF (taken to be 10%) of WDM equipment whose efficiency does not improve over time [1], [3]. Thus, the estimated power rating of an item of WDM equipment at a future point in time can be estimated given a reference power rating $\frac{P_0}{C_0}$ at year $t = 0$ and the yearly improvement rate α (e.g. 10%) as follows:

$$\frac{P}{C}(t) = \frac{P_0}{C_0} \left[OF + (1 - OF)(1 - \alpha)^{t-t_0} \right] \quad (4)$$

Sleep mode and power-aware routing — A popular research topic to improve energy efficiency is the usage of sleep modes, with examples of such work including [27], [28],




Fig. 2. Power scaling with load. (a) typical links, (b) perfectly power proportional links

[29], [30]. The general idea is that networks operate at low average utilization due to overprovisioning. However, power consumption remains almost independent of the actual load [30]. This is illustrated in Fig. 2(a). As such, shutting down equipment when it is not in use could lead to substantial overall savings. We have already discussed in Section III-B the application of sleep modes on backup links of protected traffic. However, communications equipment at low-utilization does not imply that links are *not* loaded. To be able to put them into a low-power sleep mode, traffic from lightly loaded links could be moved to other links with spare capacity. As a result, the rerouted traffic will likely not be routed over the shortest path anymore; but under the correct constraints could still result in energy savings. Examples of such power-aware routing can be found in [30], [31], [32], [33]. One final aspect to consider with respect to sleep modes is the timing associated with powering-up and -down of (sub)equipment. These timings can be unreasonably long with respect to required response times. Unfortunately, to our knowledge little to no public information is available to give an indication of realistic timings.

More power proportional systems — A more general approach to sleep mode, is to make equipment more power proportional. Power proportionality implies that the power scales linearly with the load, and is depicted in Fig. 2(b). Power proportionality applied to computing is discussed at length in [34], but also applies equally well to network equipment, as well as to line cards, and also to complete shelves and racks. Note that, as the power rating value in our model is based on capacity (i.e. maximum load), and not the actual load, the influence of varying traffic load (in its current form) is not captured in the power rating factor. Relatedly, in [10] we have looked into the impact of equipment filling levels on the power rating value. This is because 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 value towards the end of the equipment’s life, when the chassis overhead is shared by the maximum number of functional components. Our results 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 to near to double its value. This makes a case for more power proportionality at a system level. Equipment deployed in racks should scale better with the actual deployed capacity, leading to a more power proportional system. This could be achieved by designing systems such that more power is consumed in the incremental parts (e.g. line cards), and less power in the common (shared) parts, e.g. typically the chassis.

E. 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], [4].

We consider the technique of optical bypass to reduce the hop count.

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 [4], [35]. The idea is that traffic not intended for the IP node, remains in the optical layer and thus bypasses the IP. The light path is switched, using ROADMs or 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 or Optical Add/Drop Multiplexer (OADM) since multiple wavelengths are switched at the same time). From an energy (and also cost) perspective, optical bypass requires adequately filled optical channels. Depending on the ratio of IP switching power rating versus transponder power rating, there will be a channel filling ratio below which optical bypass consumes more energy than IP switching and grooming. This is because each demand using optical bypass requires a *dedicated* optical channel, and associated transponder pair. If, in contrast, traffic is pulled up to the IP layer in each node, demands on the same outgoing link can be groomed (i.e. bundled), which saves on the transponder pair’s power consumption.

Note that some approaches to reduce overall network power consumption can actually result in an *increase* in the hop count. This is the case with a number of approaches exploiting sleep modes, as discussed in Section III-D.

IV. CONCLUSION

In this paper we have surveyed a number of approaches to reduce the energy consumption of backbone telecommunication networks. While making a list of potential power reduction techniques has been done before, our approach has been based on mapping the techniques onto an analytical power consumption model. By doing so, a more intuitive

understanding of the impact of specific approaches, along with their relative importances and relationships to each other can be obtained. It is clear there are a lot of potential techniques to reduce power consumption in telecommunication (backbone) networks. However, since the power reductions achievable in the external overhead, traffic protection and hop count are inherently limited to $2\times$, $2\times$ and $3\times$ to $4\times$ respectively, it is clear that achieving backbone power savings larger than 12 to 16 times can only be realized through techniques that reduce the amount of traffic and the power ratings of the telecoms equipment. To actually quantify the possibilities for partial and total power savings, a more quantitative study now needs to be performed for each of the techniques discussed in this paper, followed by a mapping onto the detailed analytical model. This is the subject of future work.

ACKNOWLEDGMENTS

The work described in this paper was carried out with the support of the IBBT-project GreenICT, the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257740 (Network of Excellence TREND) and grant agreement n. 247674 (STRONGEST project).

REFERENCES

- [1] 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.
- [2] Cisco, "Cisco Visual Networking Index: Forecast and Methodology, 2010-2015," 2011.
- [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.
- [4] 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.
- [5] M. Parker and S. Walker, "Roadmapping ICT: An Absolute Energy Efficiency Metric," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 3, no. 8, pp. A49–A58, 2011.
- [6] J. Restrepo, C. Gruber, and C. Machuca, "Energy profile aware routing," in *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*, pp. 1–5, Ieee, 2009.
- [7] 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.
- [8] 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*, no. 99, pp. 1–22, 2010.
- [9] W. Van Heddeghem and F. Idzikowski, "Equipment power consumption in optical multilayer networks - source data," Tech. Rep. IBCN-12-001-01, IBCN, University of Ghent, January 2012. Available at: <http://powerlib.intec.ugent.be>.
- [10] 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 *submitted to IEEE Online conference on green communications (GreenCom)*, 2012.
- [11] EPA, "Report to Congress on Server and Data Center Energy Efficiency," tech. rep., U.S. Environmental Protection Agency, 2007.
- [12] "Data center efficiency - power usage effectiveness." <http://www.google.com/about/datacenters/inside/efficiency/power-usage.html>, 2012.
- [13] 42u, "What is PUE? What is DCiE?." <http://www.42u.com/measurement/pue-dcie.htm#What-is-PUE>
- [14] N. Rasmussen, *Implementing energy efficient data centers (white paper 114, revision 1)*. APC, 2011.
- [15] 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.
- [16] "Data center best practices." <http://www.google.com/about/datacenters/best-practices.html>, 2012.
- [17] 42u, "Cooling solutions for server racks, computer rooms, server rooms and data centers." <http://www.42u.com/42u-rack-cooling.htm>.
- [18] 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.
- [19] 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, August*, vol. 3, pp. 76–87, 2006.
- [20] X. Fan, W. Weber, and L. Barroso, "Power provisioning for a warehouse-sized computer," *ACM SIGARCH Computer Architecture News*, vol. 35, no. 2, pp. 13–23, 2007.
- [21] J. Vasseur, M. Pickavet, and P. Demeester, *Network recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*. Morgan Kaufmann Publishers, 2004.
- [22] 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.
- [23] J. Kooimey, 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, no. 3, pp. 46–54, 2011.
- [24] T. Ebrahimi, "What future for image/video compression." http://www.icme2011.org/files/icme2011_panel1_ebrahimi.pdf, July 2011.
- [25] G. Peng, "CDN: Content Distribution Network," *Arxiv preprint cs/0411069*, 2004.
- [26] 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*, pp. 1–5, IEEE, 2011.
- [27] S. Nedeveschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing network energy consumption via sleeping and rate-adaptation," in *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, pp. 323–336, 2008.
- [28] L. Chiaraviglio, M. Mellia, and F. Neri, "Reducing power consumption in backbone networks," in *Communications, 2009. ICC'09. IEEE International Conference on*, pp. 1–6, IEEE, 2009.
- [29] 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.
- [30] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," in *Proc. of INFOCOM, Phoenix, USA*, pp. 1130–1138, April 2008.
- [31] P. Wiatr, P. Monti, and L. Wosinska, "Green lightpath provisioning in transparent WDM networks: Pros and cons," in *2010 IEEE 4th International Symposium on Advanced Networks and Telecommunication Systems*, IEEE, Dec. 2010.
- [32] J. C. Cardona Restrepo, C. G. Gruber, and C. Mas Machuca, "Energy Profile Aware Routing," *2009 IEEE International Conference on Communications Workshops*, pp. 1–5, June 2009.
- [33] B. Puype, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, "Multilayer traffic engineering for energy efficiency," *Photonic Network Communications*, vol. 21, no. 2, pp. 127–140, 2011.
- [34] L. Barroso and U. Holzle, "The case for energy-proportional computing," *Computer*, vol. 40, pp. 33–37, December 2007.
- [35] G. Shen and R. Tucker, "Energy-minimized design for IP over WDM networks," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 1, no. 1, pp. 176–186, 2009.