Energy-Efficiency in Telecommunications Networks: Link-by-Link versus End-to-End Grooming

Ward Van Heddeghem, Maarten De Groote, Willem Vereecken,
Didier Colle, Mario Pickavet and Piet Demeester[1]

Abstract—The large share of energy consumption in telecommunication networks is expected to shift from access networks to core networks. Estimating the power consumption of core networks is not easy, as they vary a lot in size and topology. Using an exemplary but realistic core network, we estimate its power consumption for both a link-by-link grooming and an optical end-to-end grooming scenario. We show that optical end-to-end grooming consumes about half the power of the alternative scenario.

Index Terms—power consumption, network concept, optical bypass, grooming

I. INTRODUCTION

The energy footprint of Information and Communication Technology (ICT) equipment has become an issue both from an economical as well as an environmental point of view [1]. However, due to the increasing diversity of ICT equipment and its high-penetration throughout society it has become no easy task to estimate the energy footprint of ICT. Determining the power consumption of ICT network equipment is particularly complex since it entangles the whole earth and boundaries are not always clear.

Generalized, the ICT network can be split up in a high number of *access networks* which are connected to each other through a smaller number of *core networks* (see Fig. 1). The access network allows end devices (such as home routers, mobile or fixed phones, PDAs, laptops, digital TVs, etc.) to connect to a core network. The core network provides a high-speed intermediate connection system that links both the access networks of the engaging end-devices. Because of its geographical span, and because several players operate in this field, multiple core networks exist, all linked together to what could be considered a single super core network.

While the access networks currently consume by far the highest share of the total energy needed by the telecommunication networks, with rising traffic volume this is expected to shift to the core networks [2]. Therefore, in this paper, we will focus on estimating the power consumption of such a core network. Since there is no default core network architecture, and since core networks range from country-sized to continent-sized, establishing a general power consumption model for core networks is not straightforward. Therefore, in this initial study we will estimate the power consumption of an exemplary core network.

In section II, we will introduce an exemplary pan-European core network, loosely based on the Géant research network [3]. We consider two scenarios: a network with *link-by-link grooming* (i.e., nodes will convert all optical traffic to the electronic domain, and perform sub-wavelength switching), and a network with *end-to-end grooming* (i.e., an optical path will be set up between the end nodes, the bypass traffic in a node will remain in the optical domain).

We present a power estimation methodology for both scenarios in section IV.

In section V, we present and discuss the resulting estimated power consumption, as well as the scaling of both scenarios to future network demands.
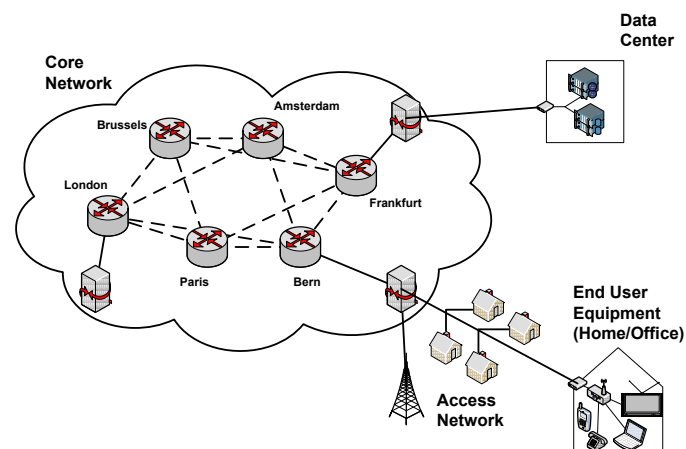

Fig. 1 – The core network situated amidst its fellow ICT equipment

[1] Manuscript submitted on December 20th, 2009. The authors are with the Department of Information Technology (INTEC) of Ghent University - IBBT, Gaston Crommenlaan 8, B-9050 Gent, Belgium (e-mail: {firstname.lastname}@intec.ugent.be).

II. RELATED WORK

In [2], Lange determines the trend of energy consumption in a telecommunication network (core, aggregation and access) based on traffic volume forecasts, and concludes that in the access network the energy consumption scales with the number of subscribers, whereas for the core network it scales with the traffic volume. As a result, with growing traffic demands the share of core networks in total energy consumption is expected rise.

In a number of papers ([4], [5] and [6]), Baliga et. al. models and estimates the power consumption (per customer) of WDM links, access, metro and core network in function of the average access rate and access technologies (such as FTTH, PON, Wimax, DSL). This is also done based on the power consumption of commercially available equipment such as routers and DSL modems. More specifically, what concerns the core network, his estimate is based on the global total switching capacity required to support a given access rate, and with a predefined share of add/drop and bypass traffic in each node. The power consumption in the core network is then a function of the oversubscription rate, connected homes and peak access rate. In another paper [4], the core network is modeled as a function of the average number of hops in the core network.

Aleksic has analyzed the power consumption in future high-capacity network nodes [7]. He considers optical and electronic architectures for both circuit switching (similar to our end-to-end grooming) and packet switching (similar to our link-by-link grooming). His study is based on a generic model of the subcomponents that make up a high-capacity network node, such as a router or switch. Summing the power consumption of these subcomponents (such as the backplane, switching fabric, line card components), he estimates the power consumption of these nodes in function of the aggregated bandwidth, and then validates his model against power consumption values of commercially available equipment where possible. He concludes that optical nodes consume generally less power than electronic ones, and that circuit switching consumes less power than packet switching.

III. CORE NETWORK ARCHITECTURE

A core network consists of a small number of high-capacity routers. These routers serve as an access point for the access network and route the traffic using high-capacity links to other, distant routers of the core network. These nodes are usually connected through high-capacity wavelength-division multiplexed (WDM) fiber links. A WDM fiber link carries multiple optical signals over one single fiber by employing wavelength division multiplexing. Each wavelength, or channel, is capable of carrying for example 10 or 40 Gbps, with 40, 80 or more wavelengths multiplexed per fiber. Typically, the WDM links in today's network comprise 10 Gb/s Packet over SONET (PoS), SDH, or 10 Gb/s Ethernet, or for longhaul links 40 Gb/s PoS or SDH [4].

A. Topology

In this paper we consider an exemplary core network as shown in Fig. 2. It is based on the pan-European Géant research network [3], but has been modified to represent a commercial transport network (for example, to protect against single link failures, the topology has been modified so that each node is at least connected to two other nodes). It contains 34 nodes, each connected through one or more WDM fiber links capable of carrying 40 channels of 10 Gbps per fiber.
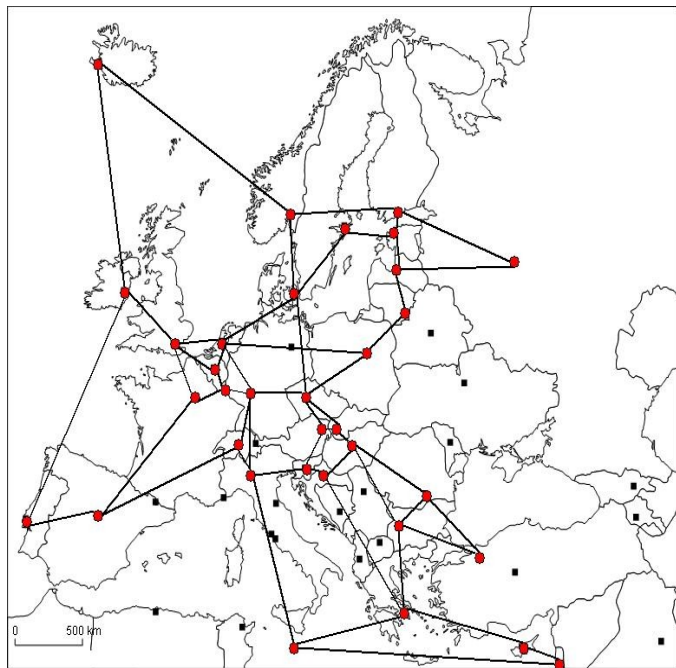

Fig. 2 – The logical topology of our exemplary core network

A traffic demand matrix (34 x 34) details the traffic demands between each node, with a smallest granularity of 1 Gbps. All traffic is bidirectional: e.g., a 4 Gbps demand from Belgium to Sweden implies a 4 Gbps demand to opposite way.

Using a shortest cycle path algorithm the traffic demands are routed through the network. The shortest cycle path algorithm provides 1+1 protection. For each node-to-node traffic demand, a data path will be set up, likely traversing multiple intermediate links and nodes. As a result, the required traffic over each link and in each node will be known.

B. Node Architecture : Scenarios

For the make-up of the core nodes, we consider two different scenarios: *link-by link grooming* and *end-to-end grooming*.

Traffic grooming is the process of grouping smaller traffic flows into larger units, such that they can be processed as a single entity at a reduced cost. This can be done for example via time division multiplexing (TDM) or wavelength division multiplexing (WDM). [8]

1) Link-by-link Grooming

In the link-by-link grooming scenario, all traffic demands on a single link are considered as one global volume of traffic. One way to implement this would be through packet switching (which can be considered as TDM). Using current technology, this can only be done feasibly in the electronic domain. Thus, in each core node *all* optical traffic on its links will need to be converted to the electronic domain, and if necessary back to optical domain (O-E-O conversion). The drawback of this approach compared to some or all traffic (if possible) remaining in the optical domain is that it requires substantially more energy [7].

In this scenario, to perform link-by-link grooming, we consider each core node to be equipped only with a *high-end core router* that matches the required traffic switching volume.

2) End-to-end Grooming

In the end-to-end grooming scenario, all traffic demands between two end nodes are considered as one global volume. This allows setting up a dedicated optical path from end node to end node, effectively setting up optical circuit switching.

To determine the equipment needed in our network nodes, we will distinguish add/drop (i.e. local) traffic, regenerated traffic and bypass traffic (see also Fig. 3).

Add/drop traffic – For each end-to-end demand an optical path is set up between the interacting core nodes. This results in so-called add/drop traffic in both nodes. Since the smallest granularity of an optical channel is 10 Gbps, the 1 Gbps demands will be converted to 10 Gbps wavelengths before the routing algorithm calculates the paths. Obviously, this will have an impact on the required core router capacity.

Bypass traffic – When both communicating nodes are not neighbors, the optical path will traverse intermediate nodes, leading to bypass traffic in these intermediate nodes.

Regenerated traffic – Even with intermediate optical amplification, the length of such an optical path is limited due to attenuation and increasing signal to noise ratio. As a result, optical paths exceeding this maximum length will require optical regeneration in an intermediate node, in effect leading to O-E-O conversion. Thus, some nodes will have a certain amount of regenerated traffic. We consider a typical length of 3000 km [9] before regeneration is necessary.

To handle add/drop traffic, we will equip each core node with a *high-end core router*. To handle regenerated traffic, we will equip the relevant core nodes with one *regenerator* per channel. Since bypass traffic remains in the optical domain, and is typically switched using a wavelength selective switch (WSS) which requires practically no energy, we can safely ignore bypass traffic in our study; it will have no relevant impact on the power consumption of a core node.
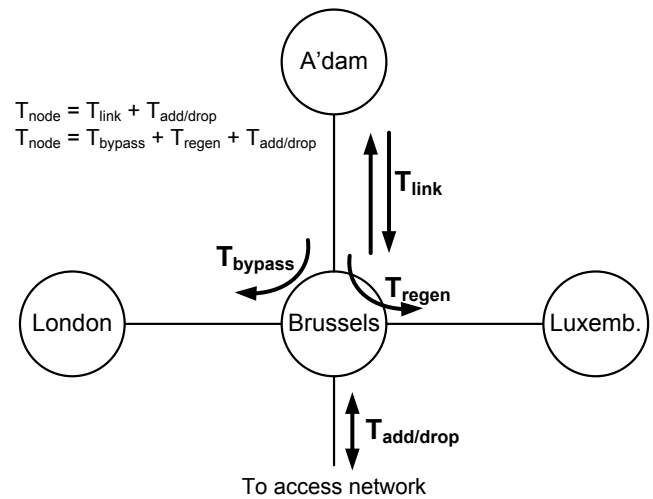

$$T_{node} = T_{link} + T_{add/drop}$$
$$T_{node} = T_{bypass} + T_{regen} + T_{add/drop}$$

Fig. 3 – Example traffic in a node, consisting of link traffic T_{link}, bypass traffic T_{bypass}, regenerated traffic T_{regen}, and add/drop traffic $T_{add/drop}$

IV. POWER CONSUMPTION METHODOLOGY

A. Core Router Power Consumption Data

As explained in the previous section, each core node will in both scenarios be equipped with a high-end core router. To calculate the resulting power requirements we will use power consumption data from commercially available high-end core routers.

The number of high-end core routers is limited, with notable manufactures including: Cisco (with the CRS-1 Router), Juniper Networks (T-series), Huawei (NetEngine 5000E) and Brocade Communication Systems (NetIron XMR series). Since we want to stick to one router type or family, and since power consumption data is readily available, we will use the T-series core router family as a reference in this study. The T-series scales from 320 Gbps (T640) up to 12.5 Gbps (TX Matrix Plus in maximum configuration) bidirectional switching capacity (see Fig. 4).
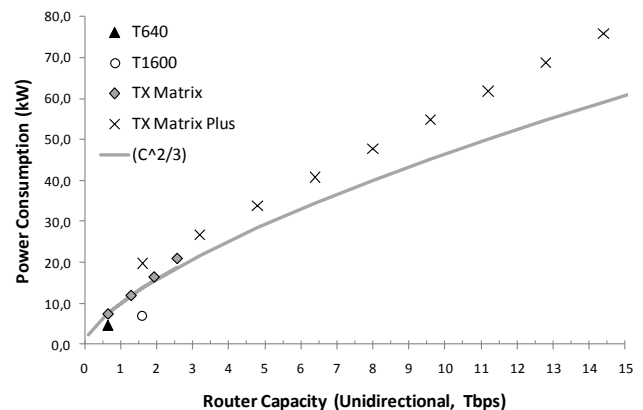

Fig. 4 – Power consumption scaling of the Juniper T-series core routers. The $P=C^{2/3}$ line shows the power versus capacity function (in W and Mbps) as proposed in [10].

Fig. 4 plots the maximum power consumption against switching capacity of the various T-series routers. The throughput values are specified as unidirectional, i.e. 1 Tbps corresponds to 500 Gbps bidirectional. The grey line indicates the power versus router capacity function (Power[W] = Capacity[Mbps]$^{2/3}$) as proposed in [10].

To translate switching capacity to power, we will use the maximum theoretical power consumption of the T-series core router with the lowest power consumption that fulfils the switching demand. For example, for a core node that has been calculated to require 1 Tbps unidirectional switching capacity, we would report the power consumption to be that of the T1600 router (7 kW), even though we can see on Fig. 4 that the 2nd TX Matrix configuration (1.28 Tbps) matches closer to 1 Tbps than the T1600 does.

The complete range of the T-series routers (T640 up to TX Matrix Plus) will be required to match the calculated node switching capacities, notwithstanding a few exceptions with nodes going below 320 Gbps required bidirectional capacity. Their influence on the general outcome of this study is marginal, and we will ignore the error introduced by using the power consumption of the T640 as a baseline.

B. Link-by-link Grooming: Node Power Consumption

To determine the power consumption in the nodes, we only need to calculate the core router power consumption. We proceed as follows.

(i)   The bidirectional 1 Gbps demands are routed using the routing algorithm described in section II. The output is a matrix of add/drop node traffic and link traffic.

(ii)   The required bidirectional switching capacity in a node is then the sum of: (a) the sum of its rounded link traffic (link traffic is rounded to 10 Gbps, the smallest granularity of line card ports in the core router), and (b) the add/drop traffic for that node, rounded to 10 Gbps.

(iii)   Using the technique detailed in section IV.A, the resulting power consumption of the router is calculated.

C. End-to-end Grooming: Node Power Consumption

To determine the power consumption in the nodes, we sum the power consumption of the core router and the regeneration equipment.

1) Core Router

To calculate the power consumption of the core router, we proceed as follows.

(i)   The bidirectional 1 Gbps demands are translated to a number of 10 Gbps wavelength channels. For demands below 10 Gbps this will result in underutilization.

(ii)   These bidirectional 10 Gbps traffic channels are then routed and assigned to nodes and links using the routing algorithm described in section II.

(iii)   The required bidirectional switching capacity for the core router in a node is then the sum of:

(a) at the core network side: the 10 Gbps add/drop traffic in the node,

(b) at the access network side: the original 1 Gbps add/drop traffic in the node (i.e., the sum of the 1 Gbps demands).

(iv)   Using the technique detailed in section IV.A, the resulting power consumption of the core router is calculated.

2) Regeneration Equipment

To calculate the power consumption of the regeneration equipment, for each node we multiply the 10 Gbps regenerated traffic (as output for each node by the routing algorithm) with the power consumption of a regeneration device.

Regeneration consists of a dedicated O-E-O conversion, thus it will consume probably less than the power consumption of two line cards. We estimate the power consumption of a regeneration device to be around 50 W.

D. WDM Link Power Consumption

Each optical fiber requires at least two optical amplifiers: one when leaving a node and a second when entering a node (see Fig. 5). In addition, to cater for signal attenuation, an optical line amplifier is required at a typical interval of about 80 km [9].
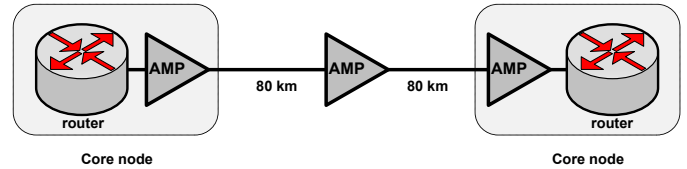

Fig. 5 – Optical amplifiers (AMP) in a WDM link: one amplifier at the begin and end, and one approximately every 80 km.

For the multi-wavelength optical amplifier of [11] a power consumption value of 6 W is reported. However, we feel this is too low a value and does not account for end-of-the-line and inline amplification monitoring support systems. Therefore, and following discussions with industrial partners, we use 25 W per amplifier.

To calculate the power consumption P_{link} for a specific link, we have to calculate the number of amplifiers per link N_A. This is a function of the link length L_{link}, and the number of fibers N_{fibers} in this link:

$$P_{link} = P_a . N_A = P_a . N_{fibers} . \left(\left(\left\lfloor \frac{L_{link}}{80km} \right\rfloor \right) + 2 \right)$$

The numbers of fiber per link, N_{fibers}, is a function of the maximum capacity of the link T_{max}, and the total link traffic T_{link} over that link. For a WDM with 40 channels at 10 Gbps we get:

$$N_{fibers} = \frac{T_{link}}{T_{max}} = \frac{T_{link}}{400Gbps}$$

The total power consumption over all the links is than the sum of all individual link power consumptions.

For both scenarios, the WDM link power consumption calculation is identical, with the sole exception that for the end-to-end grooming scenario we have to consider 10 Gbps channels.

V. RESULTS

Table 1 and Fig. 6 show the result of the power consumption estimation. The estimation also includes the results for a 50% annual traffic increase over 5 years.

A. Initial Year

For the initial year, the total power consumption of both scenarios is almost equal, around 750 kW. This compares roughly to the peak output of a medium-sized wind turbine (height 50 m, diameter 40 m) [12]. This figure already includes a power usage effectiveness (PUE) factor of 2 for the core router and regeneration [13]. Thus, half of the power (i.e., 375 kW) is consumed by the core router and/or regeneration device, the other half is consumed through overhead such as cooling equipment. The bulk of these 750 kW is in the node power consumption, with the WDM links consuming only about 4% (link-by-link grooming) and 10% (end-to-end grooming) of the total power.

The fact that both scenarios consume about the same power might be surprising at first – after all, part of the traffic in the end-to-end grooming scenario, i.e. the bypass traffic, is almost 'power consumption free'. This can easily be explained however. The routers are dimensioned to the 10 Gbps add/drop traffic on the core network side, even though these 10 Gbps are far from optimally used. To illustrate, in the traffic demand matrix, more than 90% of the demands are below 5 Gbps.

Table 1: Projected evolution of the power consumption (kW) over 5 years for both scenarios, following a traffic increase of 50% per year

	Year				
	1	**2**	**3**	**4**	**5**
Link-by-link scenario					
Nodes[*] (kW)	718	1253	1785	2340	3022
Links (kW)	29	35	43	53	69
TOTAL (kW)	**747**	**1288**	**1827**	**2393**	**3090**
Growth		72%	42%	31%	29%
Watt/Gbps[+]	156	166	156	141	121
End-to-end scenario					
Nodes[*] (kW)	718	858	996	1233	1556
.**Regeneration**[*] (kW)	103	111	125	151	197
.**Core router**[*] (kW)	615	747	871	1081	1359
Links (kW)	78	79	88	100	127
TOTAL (kW)	**796**	**937**	**1084**	**1333**	**1684**
Growth		18%	16%	23%	26%
Watt/Gbps[+]	166	121	93	78	66

[*] Includes a PUE factor of 2 to account for e.g. cooling
[+] This is calculated based on the total *demand* traffic volume

Also, and for the same underlying reason, the power consumption of the WDM links in the end-to-end grooming scenario will be more substantial (almost 10%, compared to 4% for link-by-link grooming). With 10 Gbps channels as smallest unity of transport, fiber capacity will fill up quicker and more fibers per link will be required, and hence more optical amplifiers.
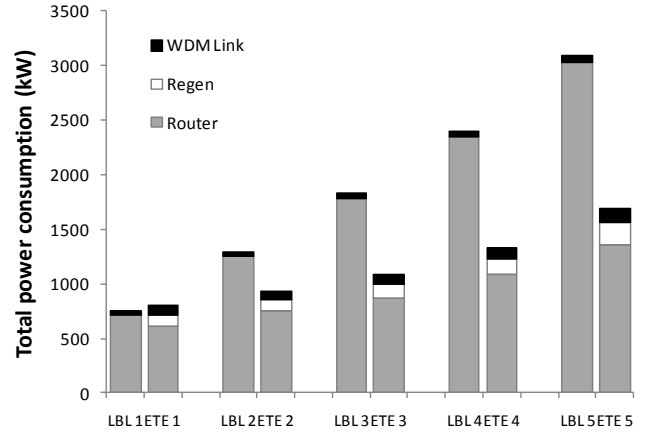

Fig. 6 – Evolution of the power consumption over 5 years for both the link-by-link (LBL) and end-to-end (ETE) grooming scenario, following a traffic increase of 50% per year

B. 5-year Projection

What happens if we apply a 50% traffic increase per year (see e.g. [14]) to the original traffic demands? From Fig. 6 we can make the following observations:

- Over time the link-by-link grooming scenario does consume less energy than the end-to-end grooming scenario. In the fifth year, this is about 50%. This is also reflected by the lower Watt per Gbps values (calculated based on the *demand* traffic volume).

- The power consumption in end-to-end grooming initially rises slower than for the link-by-link grooming scenario.
  This is because the underutilized 10 Gbps channels can carry the additional traffic demands at almost no energy increase, whereas for link-by-link this is not the case.
  This explanation is also consistent with the observed slow increase of the growth rate (from 18% to 26%): the traffic demand matches better with the 10 Gbps channels.

- Both growth rates (see Fig. 7) seem to converge to somewhere around 30%. This is in line with the power P (in W) versus router capacity C (in Mbps) relation proposed in [10]:

$$P = C^{\frac{2}{3}}$$

  For a traffic volume increase of 50% per year, this translates to 31% ($1.5^{2/3}=1.31$) increase of power per year.

- The regeneration of the optical signals accounts for approximately 12% of the total power consumption. Thus, extending the maximum reach of optical signals can potentially reduce the power consumption by about 10%.
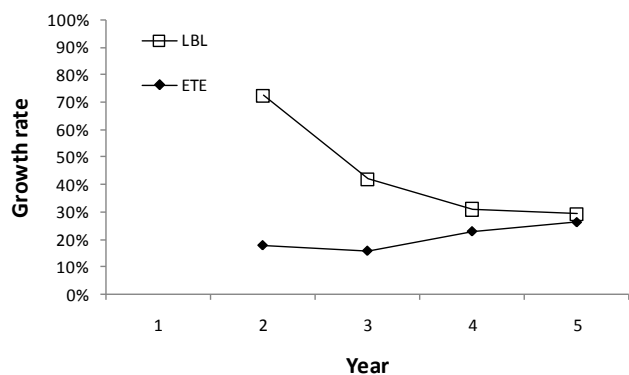
Fig. 7 – Evolution of the power consumption growth rate for both the link-by-link (LBL) and end-to-end (ETE) grooming scenario.

VI. Conclusions and Further Work

Our exemplary pan-European core network, fed with realistic traffic demands, allowed us to estimate its absolute power consumption, which is in the order of the peak capacity of a medium wind turbine. By performing optical end-to-end grooming (e.g., through circuit switching), we can save up to 50% of energy if the individual traffic demands match fairly well or exceed the capacity of the optical channels. With optical end-to-end grooming the required Watts per Gbps are about half of those for link-by-link grooming. When traffic demand increases, power consumption in end-to-end grooming does grow at roughly the same rate as for link-by-link grooming.

This is an initial study, and ample work remains to be done. Although the practical relevance of the absolute power consumption value is low, it would be interesting to benchmark this value against real-life figures. The methodology to calculate the link power consumption can potentially be simplified. The impact of the network recovery strategy on the results should be studied as well. Finally, investigation of dependence and convergence of the power consumption growth rate in both scenarios would be interesting.

VII. Acknowledgments

VIII. References

[1] M. Pickavet et al, "Worldwide Energy Needs for ICT: the Rise of Power-Aware Networking," *IEEE ANTS 2008*, Bombay (India): 2008.

[2] C. Lange, D. Kosiankowski, C. Gerlach, F. Westphal, and A. Gladisch, "Energy Consumption of Telecommunication Networks.," *ECOC, 35th European Conference on Optical Communication*, Vienna (Austria): 2009, pp. 1-13.

[3] The Géant network, http://www.geant.net/Network/NetworkTopology

[4] 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, 2009, pp. 2391-2403.

[5] J. Baliga, K. Hinton, and R.S. Tucker, "Energy Consumption of the Internet," *Joint Int. Conf. on Optical Internet, 2007 and the 2007 32nd Australian Conf. on Optical Fibre Technology. COIN-ACOFT 2007*, Melbourne (Australia): 2007.

[6] J. Baliga, R. Ayre, W.V. Sorin, K. Hinton, and R.S. Tucker, "Energy Consumption in Access Networks," *Optical Fiber Communication Conference and Exposition and The National Fiber Optic Engineers Conference*, San Diego, California, USA: 2008.

[7] S. Aleksic, "Analysis of Power Consumption in Future High-Capacity Network Nodes," *Journal of Optical Communications and Networking (JOCN)*, vol. 1, 2009, pp. 245-258.

[8] Wikipedia: http://en.wikipedia.org/wiki/Traffic_grooming, retrieved 2009-12-17

[9] M. Gunkel, R. Leppla, M. Wade, A. Lord, D. Schupke, G. Lehmann, C. Furst, S. Bodamer, B. Bollenz, H. Haunstein, H. Nakajima, and J. Martensson, "A Cost Model for the WDM Layer," *2006 International Conference on Photonics in Switching*, 2006, pp. 1-6.

[10] R.S. Tucker, *Modelling Energy Consumption in IP Networks*, retrieved from: http://www.cisco.com/web/about/ac50/ac207/crc_new/events/assets/cgrs_energy_consumption_ip.pdf

[11] MRV Optical Amplifiers datasheet, www.mrv.com/datasheets/FD/PDF300/MRV-FD-OAB_HI.pdf

[12] Wikipedia: http://nl.wikipedia.org/wiki/Windenergie, retrieved 2009-12-17

[13] The Green Grid, "The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE," Technical Committee White Paper, 2008

[14] Global IP Traffic Forecast and Methodology, 2006–2011, , 2008 [Online]. Available: http://www.cisco.com

[15] A. Gladisch, C. Lange, and R. Leppla, *Power efficiency of optical versus electronic access networks*, Brussels: ECOC 2008, 2008.

[16] C. Lange and A. Gladisch, "On the Energy Consumption of FTTH Access Networks," *Optical Fiber Comminucations (OFC)*, 2009, pp. 1-3.