

Energy-efficient traffic engineering

Bart Puype, Ward Van Heddeghem, Didier Colle, Mario Pickavet, Piet Demeester

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

Abstract The energy consumption in telecommunication networks is expected to grow considerably, especially in core networks. In this chapter, optimization of energy consumption is approached from two directions. In a first study, multilayer traffic engineering is used to assign energy-efficient paths and logical topology to IP traffic. The relation with traditional capacity optimization is explained, and the multilayer traffic engineering strategy is applied for daily traffic variations. A second study considers the core network below the IP layer, giving a detailed power consumption model. Optical bypass is evaluated as a technique to achieve considerable power savings over per-hop optical-electronic-optical 3R regeneration

1 Introduction

The interest in the energy footprint of Information and Communication Technology (ICT) has only recently been sparked. One of the first papers on this topic is probably “Greening of the Internet”, written by Gupta and Singh in 2003 [1], and discusses the power consumption of the Internet and a number of approaches to increase its energy-efficiency. Since then, the number of related publications has been rising steadily and national and international projects have sprung up.

There have been three main drivers behind this increased interest from research communities, companies and governmental bodies alike. Rising energy prices start to affect both end-users and companies and provide a financial stimulus to choose energy-efficient devices. Technically, increased operating times for mobile devices, and energy supply and heat disposal issues for large systems, press to provide solutions featuring reduced power consumption. Finally, the looming global climate change scenario has increased funding for research to reduce electrical energy consumption – as a major contributor to greenhouse gas emissions – across a wide range of technologies, including ICT equipment.

If we look at the actual figures, the energy consumed by ICT equipment was estimated to be about 4% of the global primary energy consumption¹ in 2008, and projected to double in 2020 [2]. This estimate includes both energy consumed for manufacturing as well as during use of the equipment. Subdividing ICT equipment further into data centers, personal computers, network equipment (data and telecom only), TV sets and a general ‘other’ category (containing devices such as telephone handsets, printers, fax machines, gaming consoles), it appears that each of these five categories roughly consumes an equal share of electrical energy during use. As such, network equipment constitutes about 15% of the ICT electrical energy consumption.

Generalized, telecommunication networks are typically split up into a large number of access networks and a single or small number of core networks. The access network allows end-users to connect to the core network. The core network provides a high-speed intermediate connection system that links the access networks of the engaging end-users. The access networks currently take up by far the highest share of the network power consumption. However, with rising traffic demands consumption is expected to shift to the core network [3][4]. It is in this context that approaches to limit or reduce the energy consumption in core networks have to be seen.

While a detailed survey of existing techniques to increase energy efficiency in ICT networks is outside the scope of this chapter, a number of high-level optimization approaches can be considered. One general approach consists of switching off components during low traffic load situations; for example by employing dynamic topology optimization through multilayer traffic engineering (see further). With the maximum number of idle components switched off, a next step consists of reducing the load on the remaining components; in effect avoiding overdimensioning the capacity. The optical bypass techniques discussed further in this chapter is an example of such an approach. Finally, the power consumption of the individual components itself should be reduced. A more extensive discussion can be found in [17].

In the remainder of this chapter, two energy optimization studies are presented. Consider the model as shown on Fig. 10.1. A first study looks at core networks from the view-point of a layer 3 operator, i.e., managing an IP/MPLS layer networks on top of circuit-switched infrastructure (typically WDM). Often, information about lower layers will be shielded, leading to an overlay operating model – this issue will extend to knowledge about lower layer energy consumption as well. The study presents energy-efficient multilayer traffic engineering building on an IP/MPLS layer power model. As can be expected, some similarities to typical capacity and Quality of Service (QoS) optimizing traffic engineering can be seen. Some results for short-term daily traffic variations are included.

¹ Primary energy includes all energy found in nature, not only electrical energy.

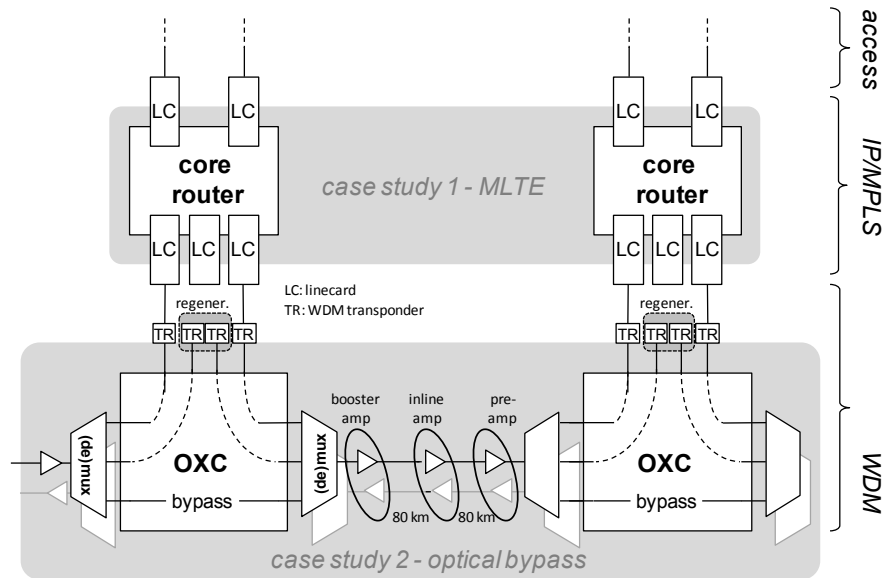


Fig. 10.1: Core network model and case study scope

On the other hand, a second study looks at core networks below the IP layer. In optical transmission systems such as WDM, the main contributing power consumers are transponders, fiber amplifiers and regeneration. Two WDM power models are presented, and in this case, IP layer energy consumption is abstracted instead. Optical bypass is evaluated as a technique to achieve considerable power savings over per-hop optical-electronic-optical 3R regeneration as the core network experiences longer term traffic volume increases.

2 Energy-efficient multilayer traffic engineering

Automatically switched multilayer IP-over-optical networks offer extensive flexibility in adapting the network to offered IP/MPLS traffic. Multilayer traffic engineering (MLTE) takes advantage of this through online IP logical topology reconfiguration in addition to the more traditional rerouting. The main functionality of MLTE is to optimize towards resource usage, bandwidth throughput and QoS performance. This functionality can be extended toward energy efficiency of IP/MPLS routing and logical topology. Moreover, these can be adapted dynamically to achieve optimal energy efficiency even for scenarios with short-term traffic changes, using the flexibility of MLTE strategies.

When including energy efficiency as an optimization goal, we need some kind of energy-based objective function or energy ‘cost’. In the next subsections, we

present a power model to abstract energy requirement allowing an IP/MPLS MLTE strategy with limited to no information about lower layers to optimize for this. Next we point out the similarities of such a model with more traditional MLTE objective functions such as capacity usage cost. Finally we adapt an MLTE strategy to include this power model and show some results comparing full mesh, slow (static) MLTE and fast MLTE under a scenario with daily traffic variations.

2.1 Power model for multilayer traffic engineering

In a general case, assuming no sub-optimal configuration, we may say that the total volume of packets processed (routed, forwarded) in a router node is dictated by the total volume of traffic offered at the interfaces. These interfaces consist of both tributary interfaces which connect to access nodes, and line cards which connect to the OXC/OADM associated with the router.

Looking at interface traffic volume, we need to make a distinction between traffic volume sent and received over tributary interfaces to the client layer, and volume carried over line cards towards the optical layer. Client layer traffic is well-defined through traffic matrices applied during simulation case studies or, in real-life situations, assumed not to be subject to operator-imposed energy limitations (which would amount to client traffic throttling). Therefore, we can say that the total amount of traffic processing in the IP/MPLS node depends on traffic volumes exchanged with its OXC/OADM, and it is these volumes that are subject to energy efficiency optimization.

Assuming all interfaces are bidirectional (i.e., connect to a RX and TX fiber), moving transit traffic from full optical switching into the electronic layer instead will increase line card bandwidth as well as router packet processing (and therefore power requirements). Of course, transit traffic in a node originates and terminates as client traffic in other nodes eventually.




Fig. 10.2: Bandwidth dependency of energy consumption for router chassis and interfaces.

Even if this transit traffic is the main area of optimization, the client traffic terminated at the node cannot simply be ignored in the bandwidth and power requirement of the core interfaces of that node. This is because energy consumption is not linear with traffic volume. This is the case both for the complete IP/MPLS router chassis and for single line cards, as shown on Fig. 10.2.

Firstly, the layer 3 router hardware itself will have a certain idle power specified. This is the power requirement when no (external) traffic is processed. Energy consumption increases with bandwidth, but there may not be a one-to-one relation, since, in addition to bulk bandwidth, traffic characteristics such as packet size or type (IP or MPLS) may also influence power requirements [12].

Secondly, ignoring traffic processing, layer 2 interfaces may show similar behavior, as shown on the right side of the figure (using logarithmic axes). For example, for an Ethernet interface, several line rates may be supported. For each line rate regime, power requirement may be unrelated to transported traffic. However, higher line rates are generally more efficient than lower ones (in Watt per Mbit/s). For example, a 1 Gbit/s line rate consumes less power than ten times the 100 Mbit/s power figure. For (electronic-optical) core network interfaces, we will typically see only 1 Gbit/s vs. 10 Gbit/s line rate, but one can assume similar effects for SONET/SDH interfaces supporting several framing rates (e.g., OC-48 vs. OC-192). Functionality such as the ability to power down parts of interface silicon in-between frames sent/received or throttling the interface to a lower line rate during lower bandwidth usage periods may lead to a smoother, more continuous curve for the interface power vs. bandwidth characteristics as well.

With the assumptions that router idle power is constant and interfaces carrying no traffic can be shut down, traffic processing is dictated by interface traffic volume. We adopt the line card power vs. bandwidth characteristic as a combined model of the bandwidth dependency of both chassis and interfaces, and assume it to be sufficient in describing power requirement optimization. In a generalized case, such a power characteristic will take the power law form $p = p_0 + (1 - p_0) b^a$ (p normalized power, b normalized bandwidth and parameters p_0 idle power, $a > 0$). Here p_0 is an abstraction of combined chassis idle power, energy dissipation in unloaded line card, etc. Simulations have shown in fact that a linear power characteristic will suffice in attaining near-optimal energy-efficient MLTE. Linear characteristics will correspond with $a = 1$; p_0 can be modified slightly to take the linear approximation into account.

2.2 Comparison with traditional capacity usage optimization

The introduction of the power vs. bandwidth characteristic suggests similarities with more traditional MLTE optimization objectives. Logical topology design and routing in traffic grooming and multilayer traffic engineering is generally centered on optimizing network resource usage for a certain traffic scenario; reducing optical layer capacity requirements, or maximizing IP layer throughput. For resource usage, we aim at minimizing number of utilized IP/MPLS links, wavelength channels, etc. For example, a resource usage performance metric could be total number of set up IP/MPLS links (as opposed to total power requirements of the IP/MPLS links).

Reducing the number of logical topology lightpaths as per traditional resource usage studies will, of course, also lower energy requirements, but the relation between logical topology and energy consumption is in fact more complex. For example, sparser logical topology meshes require longer multi-hop IP paths, which lead to increased transit traffic processing energy consumption, as explained in the previous section.

Clearly, energy-efficient logical topology design is neither about maximum optical capacity saving (sparse mesh) nor maximum IP/MPLS router power reduction (short paths; full mesh). Optimal energy efficiency should be reached somewhere in between end-to-end and point-to-point grooming. However, link power consumption p vs. traffic load b is a non-linear function; the shape of this function will determine the optimal parameters for the MLTE algorithm.

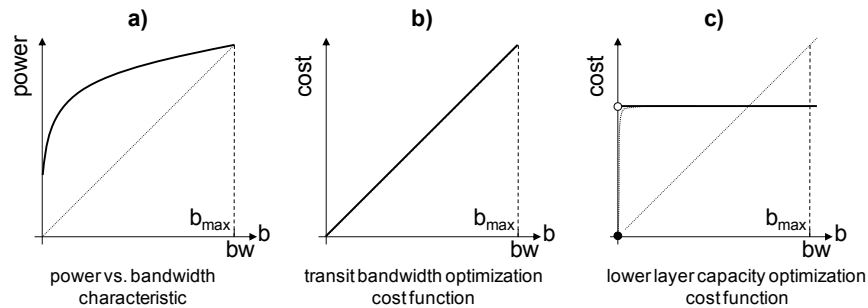


Fig. 10.3: Bandwidth dependent cost functions used in optimization.

A sample power vs. used link bandwidth function is shown on Fig. 10.3(a). When it is adopted as an MLTE cost function, the objective achieved is the minimization of total power consumption, which is the sum of all logical link $power(b)$ values. Used as a cost function, it relates to other basic optimization scenarios with their own distinct ‘cost’ shape:

- *Minimization of IP/MPLS router traffic volume, Fig. 10.3(b)* – Router traffic volume is directly related to link traffic volume; therefore such optimization

corresponds with a proportional cost (vs. link bandwidth). Lightpaths in this case are ‘free’ according to the cost function (meaning there is no optimization towards optical capacity requirements), therefore an unloaded link has zero cost. Such optimization will lead to a (almost) full mesh logical topology, with minimal bandwidth utilization in all IP/MPLS links.

- *Minimization of optical layer capacity, Fig. 10.3(c)* – In this case IP/MPLS router transit processing capacity and very long IP/MPLS paths are assumed to be ‘free’. This suggests a constant cost regardless of bandwidth utilization; except for unloaded links, which can be assigned a zero cost (they can be removed from the logical topology for an optical layer capacity reduction). Using this cost function, the objective is to minimize the total number of IP links, i.e., lightpaths.

Of course, a typical MLTE capacity optimization scenario combines the two fringe cases. Additionally, one can assume that for multi-hop grooming (as is typical with MLTE) to appear in the network, we will require a concave shape in this link cost vs. bandwidth utilization function, or at least a cost function which lies higher than the proportional one for $b < b_{max}$. Note that the function on the right is not constant over the interval $[0, b_{max}]$ as one may think at first glance but in fact (discontinuously) concave because of the zero cost for unloaded ($b = 0$) lightpaths. The middle function is neither concave nor convex, indeed there is no reason to groom traffic on multi-hop paths with such a function, the resulting routing will consequently be a shortest-path one on an IP/MPLS full mesh logical topology.

2.3 Multilayer traffic engineering strategy

The multilayer traffic engineering algorithm [13] adopted in this section uses a routing cost function (not to be confused with the power characteristic which is used to calculate power requirement after an MLTE solution is found) as shown on Fig. 10.4.

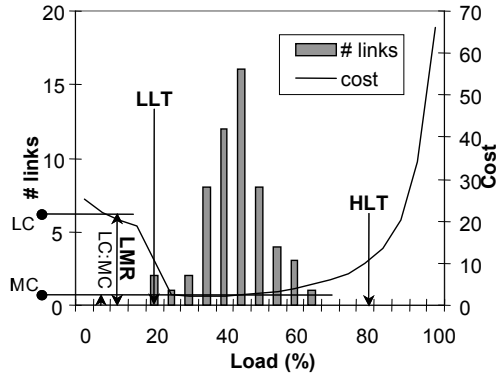


Fig. 10.4: IP/MPLS MLTE routing cost function.

The routing cost function is dependent on IP/MPLS link load, and is such that an optimal cost is seen for links between a certain range LLT – HLT (Low and High Load Threshold respectively). High costs beyond HLT avoid overloading IP/MPLS links; a cost ‘bump’ is added in order to raise costs for links with load below LLT. This makes sure that no lightly loaded links are present in the network. In fact, when routing separate traffic streams, any IP/MPLS links that have their load drop below LLT will eventually become more and more costly, leading to all traffic being deviated from those links. When applying such an algorithm on a full mesh, this leads to some of the links not carrying any traffic, with many streams being routed along multi-hop paths on the remaining links. The algorithm can route over a virtual full mesh first; the subset of links that do carry traffic after the MLTE finishes is then matched on the actual multilayer IP-over-optical network, using flexible optical switching and lightpath setup. Links not carrying traffic do not need setup, which is how the strategy optimizes logical topology (as well as routing).

Eventually, when looking at how IP/MPLS are then loaded in the network, we can plot a load histogram. An example histogram is also shown on Fig. 10.4, showing number of IP/MPLS links with a certain load for a sample traffic pattern, using this routing cost function. Obviously, the link load distribution is important as the power characteristic depends on core interface bandwidth utilization, i.e., IP/MPLS link load. This distribution is not only influenced by MLTE parameters such as the LLT, but also by the total traffic offered to the network, or rather how total traffic relates to link capacity b_{max} . For example, for low traffic volumes a very sparse (tree-like) logical topology may suffice, but this will not be the case for higher volumes.

2.4 Energy-efficient MLTE for daily traffic variations

In this section, we look at typical operator networks, and note that often such networks see daily traffic variations. Although MLTE is suited for very large traffic variations, there is still merit in adapting the network to the typical daily alternation between peak and off-peak volumes seen in access as well as core networks. In [14], the rate of off-peak vs. peak volume is 50%. Depending on the type of network and devices accessing it, this rate may be 20% for WiFi networks, or even 10% for a network servicing smart phones [15]. Obviously, a large part of the background off-peak volume is always-on traffic such as P2P services. However, even on wired access networks where such services are prevalent, daily variations are increasing in recent years [14]. Peak (day) traffic largely consists of more interactive services, which need increasingly larger bandwidths (e.g., consider gaining popularity of on-demand streaming video), making for a stronger diurnal traffic oscillation.

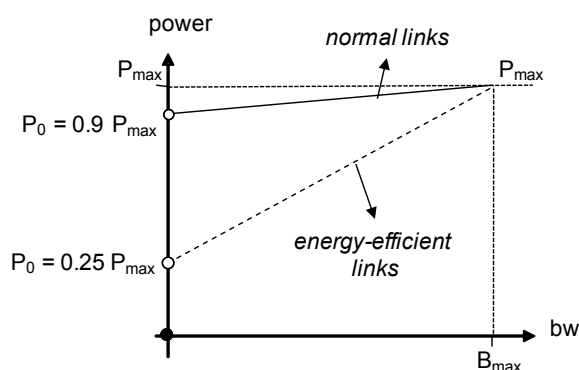


Fig. 10.5: Power characteristic for normal and energy-efficient links.

As MLTE provides an optimized topology adapted to offered traffic demand, we assume that power requirement can be optimized accordingly for the special case of daily variations, switching links on and off along with traffic volume changes. There clearly is room for specific energy efficiency-oriented optimization by somehow modifying the MLTE strategy. For the case of optical layer resource usage, we used a so-called optical metric in [16]. This optical metric provides a cost figure for each possible lightpath, each corresponding to an IP/MPLS link, so the optical metric leads to a cost assignment in the IP/MPLS layer as well. It is incorporated into the MLTE strategy by multiplying it with the original MLTE routing cost function and using this resulting function instead for routing. The optical metric in [16] is a linear function of the number of optical hops.

In the study presented here, we repeat this approach for a power-based metric: we take the interface power characteristic (as shown on Fig. 10.5) and multiply it with the original routing cost function (Fig. 10.4) to reach a new cost function

yielding automatic energy efficiency optimization. This new cost function incorporates a power metric (similar to the optical metric from [16]) that allows to take energy efficiency of IP/MPLS links into account. Specifically, it allows preferring those energy-efficient links over the less efficient links in routing the client traffic over the network and constructing the logical topology. The distinction between normal and energy efficient links is determined by the idle power requirement of the link interfaces, as shown on Fig. 10.5 for two types of links: normal ones which have an idle power P_0 of 90% of P_{\max} , and energy efficient ones at 25% P_{\max} .

To indicate the impact of MLTE optimization with a power-based metric and how it interacts with the presence of links of different energy efficiency in a network, we consider an example scenario. The sample network consists of 14 nodes, of which four have power-efficient line cards ($P_0 = 25\%$ of P_{\max} as on Fig. 10.5). For the peak volume traffic pattern, source-destination traffic is uniformly distributed between 0 and 50% maximum link capacity B_{\max} . To evaluate savings under daily traffic variations, we set the off-peak (i.e., night) traffic volume at 25% of the peak demand.

In a first case, we look at the impact of the MLTE optimization itself, by comparing it to full mesh routing. Furthermore, we run the MLTE optimization in any of two regimes, which have different reaction times. On Table 10.1, total network power requirement P_{tot} after routing (normalized as $P_{\text{tot}} / P_{\max}$) is shown for both peak and off-peak demand, and this for full mesh routing and the two MLTE regimes.

The full mesh routing requires an IP/MPLS link to be set up between any two of the 14 IP/MPLS nodes; we notice some small reduction in power requirement for the off-peak volume, but since the idle power is quite substantial for most line cards, the savings are minimal (15%). It does however set a benchmark figure for the difference between peak and off-peak power requirements under MLTE optimization techniques.

Both MLTE regimes use the same power-based metric, but have different reaction times. The first one, ‘slow’, has logical topology updates slower than the daily variations, i.e., a reaction time in terms of days. This may in fact even be some kind of dynamic grooming which acts on a large time-scale (months). In any case, the general idea of such optimization is to cope with slowly increasing traffic demands over longer time periods, as seen in most networks.

For the ‘fast’ variant, we allow MLTE actions in-between peak and off-hour periods, but this requires protocol support (e.g., GMPLS) for fast resource reservation and lightpath setup/teardown. The ‘slow’ variant necessarily optimizes for peak volume (so there is always enough bandwidth available for carrying traffic), which is why power requirements are equal for ‘slow’ and ‘fast’ during peak hours. Still, just the ability to optimize logical topology and switch off some of the links that are required for a full mesh routing, leads to a power reduction of 30% compared to the full mesh peak case.

For off-peak hours, power requirements are 45% lower (vs. full mesh off-peak) for ‘slow’ MLTE simply because of lower traffic volume, similar to the full mesh case). More importantly, for ‘fast’ MLTE, the MLTE optimization yields additional energy savings: drop of 70% vs. full mesh off-peak, since the ‘fast’ reoptimization switches off unneeded links and reroutes traffic accordingly as traffic volume drops to off-peak numbers.

Table 10.1: Peak vs. off-peak power requirements ($P_{\text{tot}} / P_{\text{max}}$) for three TE scenarios.

TE scenario	peak (100% traffic volume)	off-peak (25% traffic volume)
Full mesh	56.4	47.7
‘Slow’ MLTE	40.0	28.0
‘Fast’ MLTE	40.0	13.2

Next, we will ignore the full mesh scenario and instead look again at the ‘slow’ and ‘fast’ MLTE cases, but in this case we will incorporate the power-based metric in the MLTE strategy. The power metric allows energy consumption optimization by taking into account the energy efficiency of links, and configures optimal traffic loads for each type of line card.

On Fig. 10.6, for the same network and traffic scenario, we show peak, off-peak (‘slow’ MLTE) and again off-peak (‘fast’ MLTE) power requirements, and this for both regular MLTE (top pie-charts) and MLTE with a power-based metric (bottom). For each case, total network power requirement is indicated relatively through pie-chart area; the darker/lighter part of each pie-chart shows total power required by all line cards having high/low P_0 respectively.

In the ‘slow’ MLTE scenario, power requirement of high P_0 link stays mostly the same during off-peak hours, as high P_0 suggests link power requirement to be mostly insensitive to a traffic volume reduction. In energy efficient links however, power requirement will lower with traffic volume; this explains why the relative portion of high P_0 (vs. low P_0) power requirement increases for off-peak traffic volume.

The power metric enabled scheme tends to avoid and shut down inefficient line cards (links), especially at off-peak (‘low’) volumes. Although the slow regime offers no reconfiguration at off-peak volumes, avoiding inefficient line cards leads to better (compared to regular MLTE) scaling of total power requirement with traffic volume.

Comparing ‘fast’ regime for regular vs. power metric, a reduction of 12% in power is seen for off-peak hours (13% for the regular vs. power metric in the ‘slow’ regime). The difference for peak traffic is minimal as no interfaces are run near their idle operating point, so differences in power efficiency are minimal also.

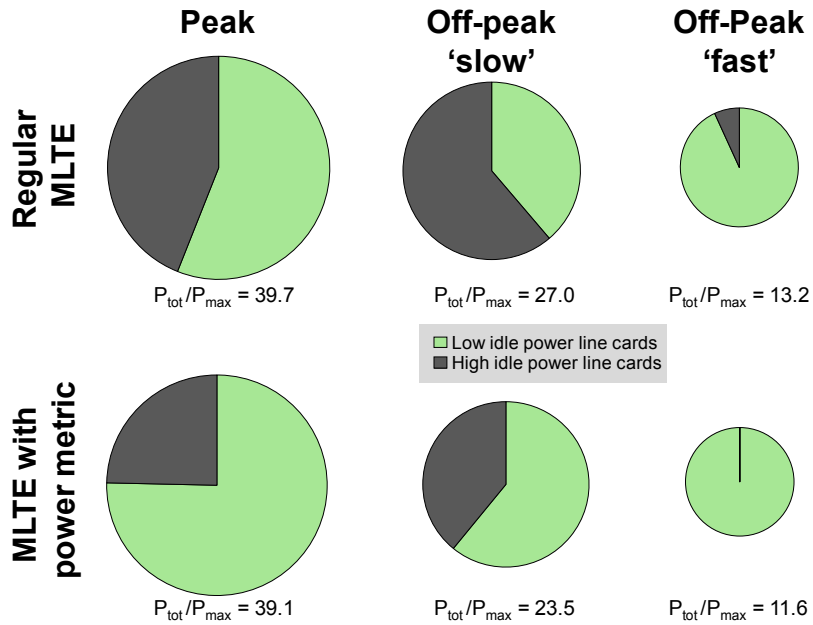


Fig. 10.6: Comparison of regular MLTE and MLTE with power metric.

Energy efficiency aware MLTE serves to traffic engineer around, and shut down power inefficient parts of the network, leading to lower total energy consumption. Even for slow response regimes, per-interface power considerations lead to logical topologies scaling better with offered traffic volume.

3 Optical bypass

3.1 Core network architecture

A core network consists of a relatively small number of core routers that are interconnected through wavelength-division multiplexed (WDM) optical fiber links, usually in a mesh or ring topology. These core routers (also referred to as layer 3 switches) serve as an access point for the access network and route the traffic using the high-capacity WDM links to other, distant routers of the core network. 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 Gbit/s or 40 Gbit/s, with 40, 80 or more wavelengths multiplexed per fiber.

Current core network architecture is typically a mix of several layers of technologies on top of each other, such as IP-over-ATM-over-SDH, as illustrated in Fig. 10.7-a. However, there is a trend to move to more homogeneous architectures where IP is routed directly over WDM links, as shown in Fig. 10.7-b.




Fig. 10.7: Evolution of the core node architecture

A technique already in use for cost reduction and router capacity offloading is optical bypass [5], shown in Fig. 10.7-c. Traffic not intended for the intermediate

node remains in the optical domain and is not processed by the router. The light path is switched, using optical add/drop multiplexers (OADMs) or optical cross connects (OXC) from an incoming fiber link directly on the appropriate outgoing fiber link. This allows reducing the capacity of the router and the corresponding power consumption. Optical bypass is possible on single-wavelength granularity or on waveband granularity (requiring fewer ports in the OXC or OADM since multiple wavelengths are switched at the same time).

As stated, while optical bypass is already in use for cost reduction and router capacity offloading, in this study we investigate the potential of single-wavelength optical bypass for power consumption reduction.

The node and link architecture supporting optical bypass, and that we consider for the rest of this paper, is shown in Fig. 10.1. The core node consists of an IP router equipped with a number of line cards (LC) having a certain capacity, e.g. 1 Gbit/s at the client side and 10 Gbit/s at the network side. The optical WDM signals are generated using long-haul transponders (TR).

3.2 Case study

For evaluating the power consumption and potential savings by employing optical bypass, we will consider the fictional but realistic pan-European core network shown in Fig. 10.8. It is based on the European Géant research network that has been modified to represent a commercial transport network; for example, each node is connected to at least two other nodes to protect against link failures.




Fig. 10.8: The logical topology of the exemplary core network

A 34-by-34 traffic demand matrix specifies the IP traffic between the nodes, with a smallest granularity of 1 Gbit/s. It represents a realistic demand scenario with randomized traffic considering higher demands for nodes located geographically nearer to each other. All traffic is assumed to be bidirectional. The demand matrix is approximately 65% meshed, i.e., 65% of all potential node-to-node demands are non-zero.

Table 10.2: Network study parameters

Parameter	Value
Network - number of nodes	34
Network - number of links	54
Network - mean node degree	3.09
Traffic matrix – mesh degree	65%
Protection	1+1
Node client-side capacity interface granularity	1 Gbit/s
Node network-side capacity interface granularity	10 Gbit/s

Furthermore, the network will provide 1+1 protection, which means that each demand is continuously transmitted over two different paths at the same time; if one path fails, the traffic will still be available without interruption over the other path.

The granularity for the interfaces available in a core node differs: the access or client-side traffic connects to a core node using 1 Gbit/s interfaces, the core network side channel interfaces are all 10 Gbit/s interfaces.

Apart from this we will consider the following two different scenarios for calculating the power consumption:

- A *non-bypass scenario*, where all the traffic in a node – both the traffic that starts or ends in the node, as well as the bypass traffic – will be processed by the IP router. This provides the opportunity for the IP router to groom – i.e., bundle – traffic from different sources destined for the same outgoing link. This assures that optical channels can be optimally filled.
- An *optical bypass scenario*, where a dedicated optical path (channel) is set up from end node to end node. This way, the bypass traffic destined for another node will not have to be handled by the IP router, and consequently not have to be converted from the optical to the electronic domain and back to the optical domain. On the other hand, if an end-to-end traffic demand is smaller than the available channel capacity, the channel will not be optimally used, resulting in a higher number of channels and equipment required.

3.3 Power consumption model

We will apply two different methods to calculate the power consumption of the total network, and then compare the resulting values:

- The first method will be based on dimensioning the network, that is, calculating for each traffic demand the path that will be followed across all nodes, and subsequently determine the equipment required. By multiplying the equipment count with its respective power consumption, the total power will be known.
- The second, alternative method is based on a simplified model or formula that takes a number of parameters and outputs the resulting total power consumption. This will be less accurate than the first method, but has the advantage of being less computationally intensive.

In both cases, we will take into account an extra amount of power consumed, mainly for cooling the locations from the heat generated by the equipment. This overhead is commonly characterized by the power usage effectiveness (PUE) [7]. The PUE is the ratio of the total amount of power consumed over the useful power consumed.

We assume a typical PUE value of 2, indicating that for each watt consumed by the actual useful equipment (such as the routers and transponder line cards), an additional watt is consumed in overhead equipment. In highly optimized and effi-

cient cooled data centers, lower PUE values are possible, but this is not yet commonplace.

3.3.1 Dimensioning-based power model

The power consumption in the core network will be the sum of the various power consuming components in the different logical layers. Descending from the IP to the optical layer, we will consider the power consumption of the *IP routers*, the *WDM transponders*, the *optical amplifiers*, and for very long optical connections the optical signal *regenerators*.

To calculate the corresponding number of each type of the above equipment, every node and link in the network has to be dimensioned. Dimensioning entails the calculation of the required number of client and network interfaces at the core nodes, and the required number of channels and fibers in the links. Following from this, we can calculate the router capacity, number of transponders, amplifiers and regenerations required. To do so, the traffic demands are routed through the network using a shortest cycle path algorithm. The shortest cycle path algorithm provides 1+1 protection. Taking all this into account, 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.

An overview of the relevant parameters and values that will be discussed in more detail in the remainder of this section is given in Table 10.3.

Table 10.3: Power consumption model values.

Element	Value Unit
Router efficiency, P_R/C_R	14 W/Gbit/s
Transponder (10G, bidirectional), P_{TR}	50 W
Optical amplifier (bidirectional, per fiber pair), P_A	100 W
Optical amplifier span	80 km
Regenerator (bidirectional, per channel), P_{RE}	20 W
Regenerator span	1500 km
Channels per fiber	40
Channel capacity	10 Gbit/s
Protection	1+1
Power usage effectiveness (PUE)	2

Router Power Consumption – The router power consumption is largely independent from the router load, as we have already mentioned in section 2 above. Therefore, the total installed router capacity is a good metric to calculate the router power consumption. The number of commercially available core routers is lim-

ited. Notable manufacturers include Cisco (CRS series), Juniper Networks (T-series) and Huawei (NetEngine 5000E).

Plotting the maximum power consumption of various router configurations of the first two vendors shows an almost linear relationship between capacity and power consumption approximated by $P_R = 14 \text{ W/Gbit/s}$ for bidirectional communication. This implies that a 1 Gbit/s interface card will consume 14 W, and a 10 Gbit/s interface card 140 W.

Transponder Power consumption – A WDM transponder is required both at the sending and the receiving node. It converts the signal from a short-range, site-specific wavelength to an ITU standardized wavelength for long haul DWDM communication, and vice-versa.

The power consumption of transponders is mainly dependent on the supported data rate. We do not differentiate by the modulation scheme or type of the transponders (for example, DPSK or coherent), as it is often an implicit requirement for achieving the required data rate.

Based on a number of vendor data sheets, power consumption values for 10 Gbit/s bidirectional transponders under typical loads are around 50 W. This value includes overhead power consumption, mainly for management of the transponders. As WDM transponders are typically only required at the network-side, we thus get for the power consumption per connection:

$$P_{transpondes}^{conn} = 2 \cdot P_{TR}$$

Amplifier Power Consumption – A number of optical amplifiers are required for each fiber pair, as can be seen in Fig. 10.1. A booster amplifier and preamplifier amplify all signals in a fiber pair respectively upon leaving and entering a node. Furthermore, an optical inline amplifier (also referred to as optical line amplifier or OLA) inserted typically every 80 km corrects for intermediate signal attenuation. Commercially available amplifiers usually can be used for all of the above three functions.

The power consumption values found in vendor data sheets vary wildly, ranging from as low as 6 W to over 600 W. However, based on a detailed analysis, a representative value for common EDFA long span amplifiers (as opposed to more powerful Raman amplifiers) is around 100 W. This value includes a 20% overhead attributed to management cards and monitoring support systems at the amplifier sites.

The total amplifier power consumption for a physical link is thus a function of the individual amplifier power consumption P_A , the number of required fiber pairs N_{fp} and the link length L_{link} :

$$P_{amps}^{link} = P_A \cdot N_{fp} \cdot \left(\left\lfloor \frac{L_{link}}{80km} \right\rfloor + 2 \right)$$

The number of required fiber pairs N_{fp} for a link can be calculated – assuming full wavelength conversion – based on the number of bidirectional channels required N_c and the number of bidirectional channel pairs N_{cfp} available per fiber pair:

$$N_{fp} = \left\lceil \frac{N_c}{N_{cfp}} \right\rceil$$

Regeneration Power Consumption – Even with intermediate optical amplification, the length of an optical path is limited due to attenuation and increasing signal-to-noise ratio. As a result, optical paths exceeding this maximum transmission distance will require optical 3R regeneration (reamplification, reshaping and retiming of the data pulse). The maximum transmission distance for this study is taken to be 1500 km [6].

Regeneration has to be performed per channel (and not per fiber), and is functionally similar to two back-to-back transponders. As such its power consumption is also dependent on the data rate. Based on a number of vendor data sheets, a typical value for bidirectional 10 Gbit/s regenerators is around 20 W. Again, this includes a (roughly) 20% management overhead.

With the number of regenerations per link depending on the number of required bidirectional channels per link N_c and the length of the bidirectional connection light path L_{conn} , the regeneration power consumption per bidirectional connection then becomes:

$$P_{regen}^{conn} = P_{RE} \cdot N_c \cdot \left(\left\lfloor \frac{L_{conn}}{1500km} \right\rfloor \right)$$

3.3.2 Alternative power model

It is also possible to estimate the network power consumption without calculating the individual equipment count. Similar to what is described in [10], this alternative model is based on a number of high-level parameters and properties of the network, and provides a way to calculate the power consumption of the same components as considered above.

The total power consumption of the core network is again simply the sum of the power consumption of the constituting components.

$$P_{core} = P_{routing} + P_{transponders} + P_{amps} + P_{regens}$$

The individual power consumption is directly given by the following equations:

$$P_{routing} = \eta_{PUE} \cdot \frac{P_R}{C_R} \cdot N_C \cdot 2 \cdot D_C \cdot (1 + \eta_{pr} \cdot H)$$

$$P_{transponders} = \eta_{PUE} \cdot \frac{P_{TR}}{C_{TR}} \cdot N_C \cdot 2 \cdot D_C \cdot \eta_{pr} \cdot H$$

$$P_{amps} = \eta_{PUE} \cdot \frac{1}{f} \cdot \frac{P_A}{C_A} \cdot N_C \cdot D_C \cdot \eta_{pr} \cdot H \cdot \left(\frac{\alpha}{80km} + 2 \right)$$

$$P_{regens} = \eta_{PUE} \cdot \frac{P_{RE}}{C_{RE}} \cdot N_C \cdot D_C \cdot \eta_{pr} \cdot H \cdot \left\lfloor \frac{\alpha}{1500km} \right\rfloor$$

Table 10.4 gives an explanation of the symbols and parameter values.

Table 10.4: Symbols and values

Quantity	Symbol	Value (non-bypass)	Value (bypass)
Core router efficiency	P_R/C_R	14 W/Gbit/s	
Transponder efficiency	P_{TR}/C_{TR}	5 W/Gbit/s	
Amplifier efficiency	P_A/C_A	0.25 W/Gbit/s	
Regenerator efficiency	P_{RE}/C_{RE}	2 W/Gbit/s	
Hop count	H	3.83	1
Provisioning factor for protection	η_{pr}		2
Provisioning factor for cooling (=PUE)	η_{PUE}		2
Number of connections	N_C	367	
Average demand per connection	D_C	See further	
Fiber filling (% of used channels in fiber)	f	100%	
Average link length	α	753 km	3.83 * 753 km

These equations can be deduced naturally from multiplying the component efficiency with the number of connections and the amount of traffic over that connection. Only the last equation expressing the regeneration power consumption will be very crude since the average link length is not a very good indicator for the number of regenerations, especially since its value is of the same order of the generation span of 1500 km.

The parameter values were determined as follows:

- The core router efficiency value P_R/C_R is as given in the previous power model, the transponder efficiency P_{TR}/C_{TR} has been derived by dividing the transponder consumption by the channel capacity, and the amplifier efficiency P_A/C_A by dividing the power consumption by the total fiber capacity (40×10 Gbit/s).
- Following the network global expectation model proposed in [11], the hopcount H in a uniform network can be approximated by the following equation, with N the total number of nodes in the network and L the number of bidirectional links in the network:

$$H \cong \sqrt{\frac{N-2}{\frac{2L}{N}-1}}$$

For the non-bypass scenario, filling in $N=34$, $L=54$ gives $H=3.83$. For comparison, the actual average hopcount is around 4.

For the non-bypass scenario, we will assume the hopcount to be 1, as direct end-to-end paths will be set up.

- The number of connections is directly available from the traffic demand matrix, as well as the average demand per connection. The latter will be increasing however.
- The fiber filling is optimistically estimated to be 100%, which will be a good approximation for large demands. The average connection length is calculated directly from the network topology. Again, it could also be estimated; the network global expectation model [11] provides an approximation based on the geographic area A covered by the network $\alpha = \sqrt{A}/(\sqrt{N}-1)$. For an estimated area of 3000×3300 km², this would give $\alpha = 653$ km. As for the bypass scenario we have assumed a hop count equal to one, the average link length will have to be multiplied by the non-bypass hop count to get a reasonable representation.

3.4 Results

3.4.1 Comparison of both models

Fig. 10.9 shows the results of applying the above two models to our exemplary European core network for an increasing traffic demand. The uninterrupted lines

show the total power consumption by dimensioning the network; the dashed lines indicate the result from the alternative model.

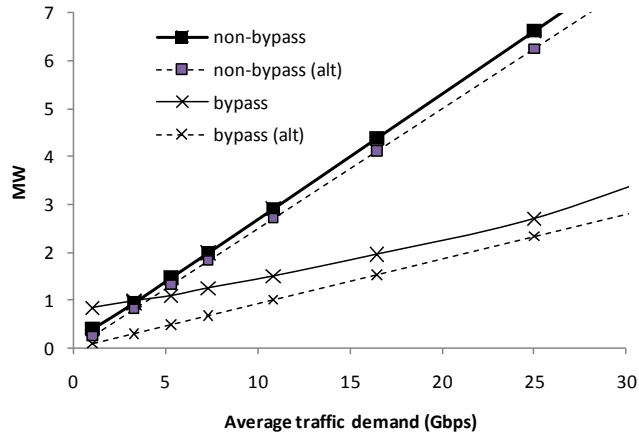


Fig. 10.9: Total power consumption of the pan-European core network as a function of the bidirectional average node traffic demand.

For the non-bypass scenario the approximation of the alternative model is reasonably accurate, for average demands higher than 10 Gbit/s (the interface) the estimate is over 93% of the actual value. A more detailed analysis shows that the deviation results from a slight underestimate of the router power consumption; the alternative model does not take into account suboptimal used interfaces.

For the bypass scenario the approximation is much less accurate, certainly for low demands. There are three main reasons why our alternative model doesn't hold well for the bypass scenario:

- at traffic demands below the network interface capacity (i.e., below 10 Gbit/s) the model does not take into account the suboptimal used interfaces, thereby overestimating the router efficiency; the underestimate is worse than for the non-bypass scenario because in the latter the grooming dampens the suboptimality,
- the amplifier power consumption is estimated much lower because a higher number of fibers will be required per link (to accommodate all the extra channels) for the same amount of demand traffic; this results in an overestimate of the amplifier efficiency,
- the regeneration estimate is too low as our simplified equation is too crude.

3.4.2 Total power consumption

As can be seen in Fig. 10.9 the total power consumption in the network is in the order of a few megawatts, scaling linearly with increasing traffic demands. Taking the non-bypass 10 Gbit/s average demand as a convenient reference, 3 MW is roughly equal to the electrical power consumption of a nearly 6000 households². From a power generation perspective, this would require about four medium to large wind turbines³.

When comparing the power consumption of both the non-bypass and bypass scenario we can make the following observations.

For very low demands (less than 3 Gbit/s) the power consumed in the bypass scenario is higher than in the non-bypass scenario. In the optical bypass scenario a dedicated optical channel is required for each end-to-end communication. As at the network side we only have 10 Gbit/s interfaces available, for demands below 10 Gbit/s the channels will not be optimally filled. This is less the case for the non-bypass scenario where all the traffic is ‘pulled’ up into the IP routing layer: the router can groom all demands for the same outgoing links, thereby optimally filling the channels, saving on the number of 10 Gbit/s interfaces and subsequently power consumption.

With rising traffic demands (from around 4 Gbit/s), the bypass strategy starts to pay off, consuming less energy. The power consumption of the bypass strategy initially rises more slowly than for the non-bypass strategy. This is because the underutilized 10 Gbit/s channels can carry the additional traffic demands at almost no energy increase, whereas for the non-bypass this is not the case. For example, consider a link which has to carry two demand streams of 4 Gbit/s. In the bypass scenario this will require two 10 Gbit/s interfaces and two channels, for a total of 20 Gbit/s. In the non-bypass scenario, the two demands will be groomed to 8 Gbit/s and will require only one 10 Gbit/s router interface. Now, if both demands double to 8 Gbit/s, the bypass scenario requirements will remain unchanged (two 10 Gbit/s, and two channels), whereas in the non-bypass scenario the groomed demands will equal 16 Gbit/s, now requiring two 10 Gbit/s interfaces and channels instead of one.

² Assuming an average power consumption of 4500 kWh/year per household, it takes about 5800 households to consume 3 MW.

³ For good measure, the average power produced by a modern wind turbine at a good site is about 30% of the peak output [8]. The peak output of a medium to large wind turbine is around 3 MW (or more), which means that in practice at least three or four wind turbines would be required.

3.4.3 Relative savings

The actual relative savings of the bypass scenario over the non-bypass scenario are shown in 10.10. As can be expected roughly when the average demand equals the channel capacity, the power savings stabilize, leveling off in this case to about 60%. Thus, by employing optical bypass we are able to consume only 40% of the power that would be used when no optical bypass would be employed.




Fig. 10.10: Relative savings of the bypass scenario power consumption over the non-bypass scenario

It is important to point out that the 60% value is no magic number. As shown in [9], the maximum energy savings achieved depend on the size of the network (in terms of nodes). For a network with similar connectivity, gains will be lower for smaller networks, and higher for larger networks. This is because for larger networks the chance of establishing longer lightpaths increases, bypassing more intermediate nodes, and thus saving on router interfaces.

3.4.4 Distribution

Fig. 10.11 shows the distribution of the power consumption among the four contributing components.

- The total power consumption is clearly dominated by the IP router, accounting for 60% to 70%.
- The transponder power consumption accounts for about 20%, with a slightly lower value for the bypass scenario: the longer end-to-end links only require a transponder at the end points, thereby skipping on the intermediate transponders in the bypass nodes.

- Amplifier power consumption is higher for the bypass scenario because the higher number of wavelengths results in more fiber pairs and thus more amplifiers.
- The regenerator power consumption is completely negligible in the non-bypass scenario where it accounts for only 0.3%. Obviously with longer optical paths in the bypass scenario, regeneration is more prominent, but still only less than 5% of the total power consumption.

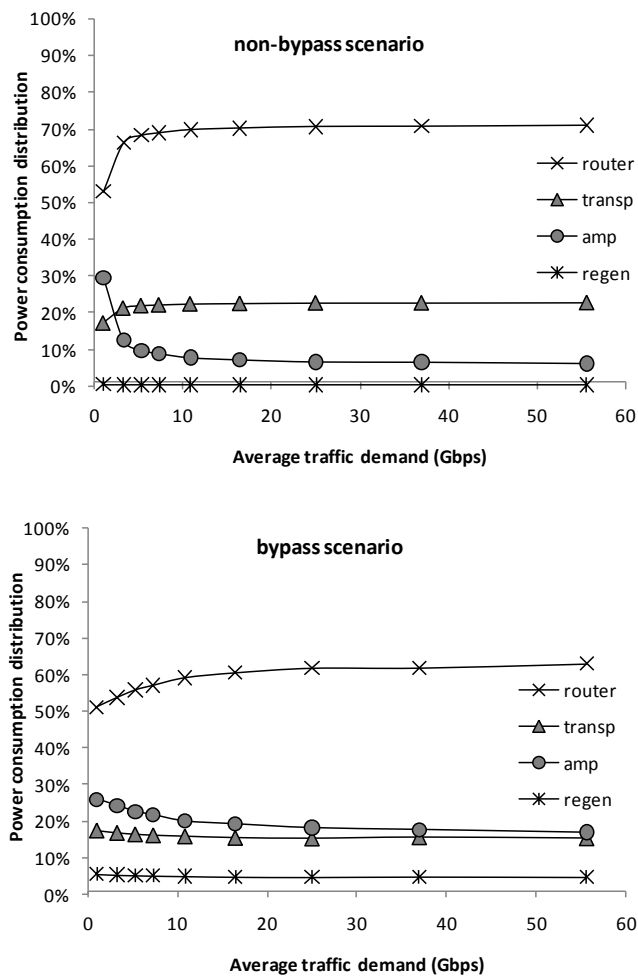


Fig. 10.11: Distribution of the power consumption

Optical bypass allows to save up to 50% of the electrical energy over a non-bypass scenario in the core network. A key requirement is sufficient filling of the channel capacity. The main share of the power is consumed by the IP routers (50% to 70% of the total power), followed by the transponders (around 20%).

4 References

- 1 M. Gupta and S. Singh, Greening of the internet, New York, New York, USA: ACM Press, 2003, doi: 10.1145/863955.863959
- 2 M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Davelder, D. Colle, B. Dhoedt, and P. Demeester, "Worldwide energy needs for ICT: The rise of power-aware networking," 2008 2nd International Symposium on Advanced Networks and Telecommunication Systems, Bombay (India): IEEE, 2008, pp. 1-3.
- 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, Jul. 2009, pp. 2391-2403, doi: 10.1109/JLT.2008.2010142
- 4 C. Lange, D. Kosiankowski, C. Gerlach, F.J. Westphal, and A. Gladisch, "Energy consumption of telecommunication networks," Optical Communication, 2009. ECOC'09. 35th European Conference on, IEEE, 2009, pp. 1-2.
- 5 J.M. Simmons and A. a M. Saleh, The value of optical bypass in reducing router size in gigabit networks, IEEE, 1999, doi: 10.1109/ICC.1999.768007
- 6 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, Oct. 2006, pp. 1-6, doi:10.1109/PS.2006.4350152.
- 7 C. Belady, A. Rawson, J. Pfleuger, T. Cader, Green grid data center power efficiency metrics: PUE and DCIE (white paper), 2008.
- 8 D. Mackay, Sustainable Energy - without the hot air, UIT Cambridge Ltd., 2009
- 9 G. Shen and R.S. Tucker, "Energy-Minimized Design for IP Over WDM Networks," Journal of Optical Communications and Networking, vol. 1, Jun. 2009, pp. 176, doi: 10.1364/JOCN.1.000176
- 10 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, 2010, pp. 1-10, doi:10.1109/JSTQE.2010.2074187
- 11 S.K. Korotky, "Network Global Expectation Model: A Statistical Formalism for Quickly Quantifying Network Needs and Costs," Journal of Lightwave Technology, vol. 22, Mar. 2004, pp. 703-722, doi:10.1109/JLT.2004.825756
- 12 J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsang, and S. Wright, "Power awareness in network design and routing," in proc. 27th Conference on Computer Communications (IEEE INFOCOM 2008), Phoenix, Arizona, April 2008, pp. 457-465. doi:10.1109/INFOCOM.2008.93
- 13 B. Puype, Q. Yan, D. Colle, S. De Maesschalck, I. Lievens, M. Pickavet, and P. Demeester, "Multi-layer traffic engineering in datacentric optical networks," in proc. the COST266/IST OPTIMIST workshop (ONDM 2003), Budapest, Hungary, February 2003, pp. 211-226
- 14 K. Cho, K. Fukuda, H. Esaki, and A. Kato, "Observing slow crustal movement in residential user traffic." in proc. 4th Annual ACM International Conference On Emerging Networking

- Experiments And Technologies (ACM CoNEXT 2008), Madrid, Spain, December 2008. doi:10.1145/1544012.1544024
- 15 M. Afanasyev, T. Chen, G.M. Voelker, and A.C. Snoeren, "Analysis of a mixed-use urban wifi network: when metropolitan becomes Neapolitan," in proc. 8th ACM SIGCOMM Conference on Internet Measurement (IMC'08), Vouliagmeni, Greece, October 2008, pp. 85–98. doi:10.1145/1452520.1452531
 - 16 B. Puype, Q. Yan, S. De Maesschalck, D. Colle, M. Pickavet, and P. Demeester, "Optical cost metrics in multi-layer traffic engineering for IP-over-optical networks," in proc. 6th International Conference on Transparent Optical Networks (ICTON '04), Wroclaw, Poland, vol. 1, 2004, pp. 75–80. doi:10.1109/ICTON.2004.1360248
 - 17 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," Communication Magazine, vol. 49, no. 6, pp. 62-69, 2011.