

Energy-aware Backbone Networks: a Case Study

Original

Energy-aware Backbone Networks: a Case Study / Chiaraviglio, Luca; Mellia, Marco; Neri, Fabio. - (2009), pp. 1-5.
(Intervento presentato al convegno First International Workshop on Green Communications (GreenComm'09) tenutosi a Dresden, Germany nel 18 June 2009) [10.1109/ICCW.2009.5208038].

Availability:

This version is available at: 11583/2285598 since:

Publisher:

IEEE

Published

DOI:10.1109/ICCW.2009.5208038

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Energy-aware Backbone Networks: a Case Study

Luca Chiaraviglio, Marco Mellia, Fabio Neri
Electronics Department, Politecnico di Torino, Torino, Italy

Abstract—Power consumption of ICT is becoming more and more a sensible problem, which is of interest for both the research community, for ISPs and for the general public. In this paper we consider a real IP backbone network and a real traffic profile. We evaluate the energy cost of running it, and, speculating on the possibility of selectively turning off spare devices whose capacity is not required to transport off-peak traffic, we show that it is possible to easily achieve more than 23% of energy saving per year, i.e., to save about 3GWh/year considering today’s power footprint of real network devices.

I. INTRODUCTION

Power consumption has become a key issue in the last few years, due to rising energy costs and serious environmental impacts of Green House Gases (GHG) emissions. Pollution and energy saving are keywords that are becoming more and more of interest to people and to governments, and the research community is also becoming more sensible towards these topics. Focusing on ICT, a number of studies estimate a power consumption related to ICT varying from 2% to 10% of the worldwide power consumption [1]. This trend is expected to increase notably in the near future. Not surprisingly, only 20% of ICT carbon emissions derive from manufacturing, while 80% arise from equipment use [2]. Moreover, among the main ICT sectors, 37% of the total ICT emissions are due to the Telecom infrastructures and devices, while data centers and user terminals are responsible for the remaining part [2].

Considering the network infrastructure alone, its power consumption accounts for an average 0.1 GW of power worldwide, as reported for example in [3]. To this extent, routers consume the large majority of energy [4], while including air conditioning and cooling can almost double the energy consumption of a network. In Italy, for example, Telecom Italia is the second largest consumer of electricity after the National Railway system [5], consuming more than 2TWh per year. It is therefore not surprising that telecom operators are trying to reduce the energy consumption of their network, although the problem is faced from a different angle.

To this extent, the study of power-saving network devices has been introduced over these years, starting from the pioneering work of [6]. In [7] some simple measurements about power consumption of networking devices are first presented; the authors consider also a network topology and evaluate the total network consumption given the power footprint of each element. Similarly, in [8] we faced the problem of defining which is the minimum set of routers and links that have to be used in order to support a given traffic demand. The idea is to power off links and even full routers while guaranteeing QoS constraints, such as maximum links utilization. Simple

algorithms have been presented to select which elements have to be powered off, and simple scenarios have been considered to assess the proposed heuristics and the achieved energy saving.

The main improvements that we present in this paper with respect to [8] are the following. First of all, the main intent of this paper is to evaluate the possible savings in an actual ISP network topology, rather than compare different heuristics like in [8]. We indeed consider a topology which is similar to the actual one adopted by one of the largest ISPs in Italy (which kindly provided the data for our study). Second, in this paper we estimate the power consumption of nodes and links using realistic figures that have been derived from available products [7]. We therefore propose a new algorithm which exploits nodes’ and links’ power consumption to select the set of elements that have to be turned off. Finally, in this paper we show that, while most network capacity has to be fully available during peak hours, traffic variation over time allows to improve the energy efficiency up to 34% during off peak hours.

While the results in this work show that there is a great opportunity to save energy consumption in a real network, it is also very true that today’s technology does not fully support the selective shutdown of links and nodes. Indeed, spare resources are provided by ISPs to provide a reliable service, so that additional links and nodes guarantee to recover from occasional failures. Keeping these additional resources always powered on is a clear waste of energy. Support to equipment selective shutdown must be explicitly introduced considering the network control plane and protocols. Indeed, while protocols like OSPF, IS-IS and BGP are capable of finding alternate routes in case of failure, they are not designed to support simultaneous “failures” of nodes and links. Finally, equipments themselves must be capable of quickly entering into and exiting from the low power status, e.g., a node or a link has to be quickly powered up to satisfy an increasing traffic demand (due to a failure).

II. TEST-CASE DESCRIPTION

A. Physical Topology

The topology considered in this work is similar to the actual topology of national ISPs. It follows a hierarchical design, as reported in Fig. 1, in which four levels of nodes are present: core, backbone, metro and feeder nodes. The inner level is composed by “core nodes” (Fig. 2.a), that are densely interconnected by 50 Gbps links. Core nodes are placed in four central Points-of-Presence (POPs) located in two cities. Each central POP hosts a pair of core nodes, each connected

TABLE I
POWER CONSUMPTION OF NODES

Node Type	Power [kW]	Fraction of Total Node Power
Core	10	9.46%
Backbone	3	19.03%
Metro	1	6.32%
Feeder	2	65.19%

to other core nodes by two links for failure protection. Central POPs may be also geographically far away, i.e. inter-POP links connecting nodes in the two cities can be 600 km long. A peering router is connected to two central POPs- to offer connectivity to the Internet by means of a 100 Gbps link.

At the second level, so called “backbone nodes” (Fig. 2.b) are connected to the core by 20 Gbps links. Each backbone node has a link to two central pops. Backbone nodes are located in chief POPs, spread in each large city. Notice that link length between the backbone and the core routers ranges between 50 and 500 km.

At the third level, “metro nodes” (Fig. 2.c) are present. Each metro node is dual-homed to two backbone nodes by 10 Gbps capacity links. Metro and backbone nodes are located in the same chief POP¹, and links between them are then short.

The last level of nodes is represented by the “feeders” (Fig. 2.d), that bring connectivity to the DSLAMs to which users are connected. Feeders aggregate traffic from users in the same neighborhood or small town. Each feeder is dual-homed to the closest pair of metro nodes by 10 Gbps capacity links. The length of links between feeders and metro ranges from 1 km to 50 km.

In this paper we consider a possible network composed by 372 routers: 8 core nodes, 52 backbone nodes, 52 metro nodes and 260 feeders. Links have a cardinality equal to 718.

To model the energy consumption of routers and links, we consider the requirements of real devices. Table I reports the mean power consumption for the different classes of nodes. Notice that these values do not consider air conditioning costs, which are responsible for up to 50% of the total power consumption.

The power consumption of links is modeled by a static contribution due to the optical transceivers, and by an additional term which takes into account possible (optical) regenerators. We consider that regeneration is required every 60/70km, and that the minimum capacity of a link is 10 Gbps, so that a 20 Gbps link consumes as two 10 Gbps physical parallel links. Then, the power consumption of link from router i to router j is given by:

$$P_{tot}^{ij} = (N_a^{ij} P_a + P_s^{ij}) [C^{ij}/10] \quad (1)$$

where $N_a^{ij} = \lceil L^{ij}/70 \rceil$ is the number of amplifiers needed to regenerate the signal (one every 70 km) for link from i to j of length L^{ij} , P_a is the power consumption of a single amplifier, P_s^{ij} is the static power consumption of router interfaces, C^{ij}

¹Notice that chief POPs are composed also by other elements, e.g the Network Access Servers (NASs) that allows user authentication. These devices are not considered in this work.

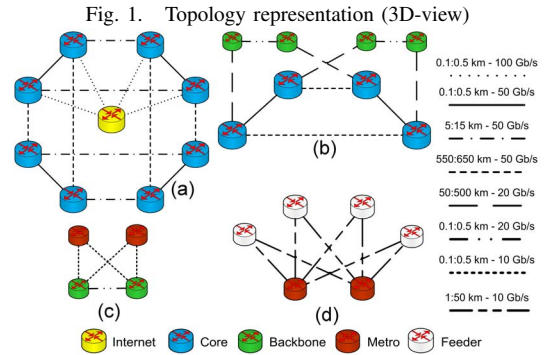



Fig. 2. Link description

is the link capacity. We assume that P_a is equal to 1 kW and P_s^{ij} is equal to 100 W.

The total power consumption of our network amounts to 1.4 MW.

B. Traffic Demand

The feeders and the Internet peering nodes are the only possible sources and destinations of traffic. Traffic estimates of the considered ISP show that about 70% of the total traffic amount is exchanged between the Internet at large and the ISP users, while the remaining part is exchanged uniformly among the feeders, i.e., 30% of traffic is confined within the same ISP, while 70% of traffic is coming from and going to other ISPs. Given $N + 1$ nodes, let t^{sd} be the average amount of traffic from node $s = 0, \dots, N$ that is going to node $d = 0, \dots, N$, i.e., $\mathcal{T} = \{t^{sd}\}$ is “traffic matrix”. Let $t^s = \sum_d t^{sd}$ and $t^d = \sum_s t^{sd}$ the total traffic generated and received by node s . Let node $i = 0$ be the peering node. Then we have:

$$\begin{aligned} t^{s0} &= 0.7t^s & \forall s \\ t^{0d} &= 0.7t^d & \forall d \\ \sum_{d=1}^N t^{sd} &= 0.3t^s & \forall s \end{aligned}$$




Fig. 3. Turning off technique

$$\sum_{s=1}^N t^{sd} = 0.3t^d \quad \forall d$$

Furthermore, we take into account QoS constraints. In particular, we assume that each link utilization cannot grow above 50% of the link capacity, so that:

$$f^{ij} \leq \frac{1}{2}C^{ij} \quad \forall i, j \quad (2)$$

where f^{ij} is the total amount of traffic flowing on link from i to j . For simplicity, we assume that t^{sd} are i.i.d. random variables, distributed according to a uniform distribution, so that $E[t^{s0}] = E[t^{0d}] = 0.7$ units of traffic, and $E[t^{sd}] = 0.3/N$ units of traffic, and $\sigma[t^{sd}] = E[t^{sd}]$.

The algorithm used to derive a possible traffic matrix works as follows: we generate a random traffic matrix $\{\hat{t}^{sd}\}$, then we route the traffic in the network according to a minimum hop path routing. In case of tie, a random path is selected among the minimum hop paths to exploit network redundancy to balance the links' load in case of multiple minimum cost paths. We then compute the amount of flow on each link, and look for the mostly loaded link $(i, j)^* = \operatorname{argmax}(\max_{(i,j)} f^{ij}/C^{ij})$. We then define a scaling factor $\alpha = \frac{C^{ij}}{2f^{ij}}$, $(i, j) = (i, j)^*$, and compute

$$t^{sd} = \alpha \hat{t}^{sd}$$

This guarantees that the constraint in Eq.(2) holds true for all links and that there is at least one link whose offered load is equal to the load bound.

III. ALGORITHM DESCRIPTION

As we detailed in [8], the problem of finding the minimum set of nodes and links that must be powered on to transport the offered traffic under QoS constraints can be formulated using an Integer Linear Programming (ILP) methodology. Unfortunately, solving the ILP is not viable, since it falls into the multi-commodity flow class, which is known to belong to the NP-hard class. Exact solutions can be found only for small networks and for some trivial cases.

In this paper, we propose an improved version of the algorithm presented in [8] that explicitly takes into account the power consumption of devices. In particular, we start with all the devices in on state; then we try to selectively power off them. The basic idea is to sort the devices according to the amount of energy they consume, and then try to power off first the devices that consume more energy. We divided the problem in two phases: first we try to power off the nodes and then we try to power off the remaining links. The generic algorithm is sketched in Fig. 3.

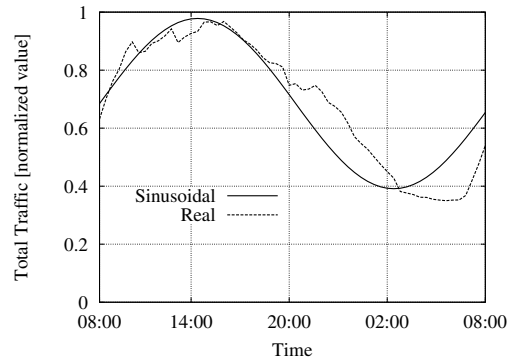


Fig. 4. Total Traffic Variation

First, we go through the ordered list of nodes and check which nodes can be powered off while guaranteeing the network connectivity and the maximum link load constraints at each step:

- 1) Sort node set in decreasing energy footprint
- 2) For each node i
 - Turn off node i and all links originating/terminating at i
 - Recompute the minimum hop paths
 - If network is disconnected, power on node i and go to next node
 - Compute all link flows by routing \mathcal{T}
 - If any link is congested then power on node i

A similar procedure is adopted for the links that are left powered on after the first step. We first sort links in decreasing order according to their power consumption. Then, we selectively try to power off them by checking if the connectivity and maximum link load constraints are met.

IV. PERFORMANCE EVALUATION

We consider a scenario in which traffic varies according to a day-night pattern. In particular, we consider both a simple sinusoidal pattern, and a real traffic profile observed on the real network. Fig. 4 reports the total offered traffic defined as $T(t) = \sum_{s,d} t^{s,d}(t)$ for both profiles; a time period of 24h is shown, values are averaged over 5min, and a normalization factor has been applied due to non-disclosure agreements with the operator. For simplicity, we assume the same traffic pattern is affecting each traffic demand, so that it can be expressed as:

$$t^{sd}(t) = f(t)t^{sd} \quad (3)$$

being $f(t)$ the shaping function at time t and t^{sd} the traffic exchanged by s and d during peak hour respectively. The shaping function has a maximum equal to 1 and a minimum which we set to 0.4 to match the minimum of the real traffic profile. Off peak traffic is then 40% of the peak-hour demand.

A. Sinusoidal Profile

In this section, we consider a sinusoidal function as shaping profile. Fig. 5 reports the node power saving, i.e., the percentage of power that is saved due to nodes that can be switched

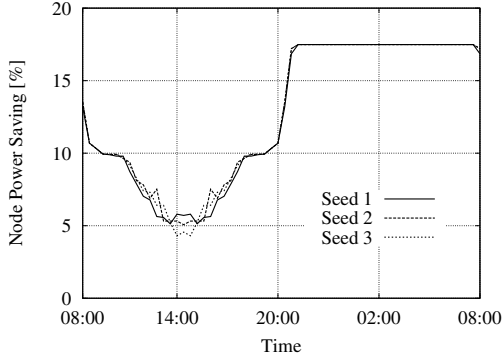


Fig. 5. Energy Saving for Nodes

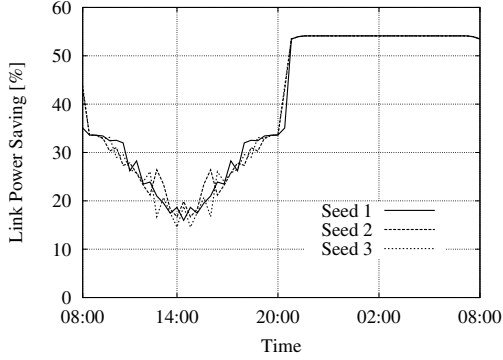


Fig. 6. Energy Saving for Links

off during a time interval. We compute the power saving as:

$$Node_{saving}(t) = \frac{\sum_i P_{on}^i(t)}{\sum_i P^i(t)} \quad (4)$$

where the numerator is the power consumed by on nodes for the energy-aware network and the denominator is the power consumed by nodes for a standard network. Three randomly generated traffic profiles are provided. Note that $Node_{saving}(t)$ is constant during night, since the connectivity is the tightest constraint, being the offered traffic much smaller than during peak hour. Saving of 18% of power is possible. As expected, during the day the node power saving decreases as the traffic increase, since more capacity is required in order to guarantee the maximum link utilization constraint. Notice however that under the considered scenario it would be possible to always turn off some nodes, so that a minimum 5% of power saving is always possible.

Fig. 6 shows the link power saving computed considering the total power consumed by links expressed by:

$$Link_{saving}(t) = \frac{\sum_{ij} P_{on}^{ij}(t)}{\sum_{ij} P^{ij}(t)} \quad (5)$$

In this case the saving is higher than in the node case since a much larger number of links can be switched off during off-peak hours. During the day instead, it is not possible to save a lot of energy. Also in this case it is possible to always turn off about 15% of links without violating the connectivity and maximum traffic load constraints.




Fig. 7. Energy Saving for each class of node

To give more insights, Fig. 7 reports the breakdown of the percentage of nodes that are switched off detailing core, backbone and metro nodes. Values have been averaged over the three different runs. The plot shows that during off-peak hours it is possible to turn off up to 50% of nodes that are not source/destination of traffic, being the backbone and metro nodes the largest fraction of them. This reflects the fact that the network has been designed to recover from possible faults, which requires additional resources. These additional resources are not exploited to carry traffic during off-peak time, and then they can be powered down to save energy. During peak hours on the contrary, the saving is much lower, as only about 10% of nodes can be powered off, being the majority of them backbone nodes. These additional nodes may be required to recover from occasional faults. A similar reasoning can be applied considering links, but it is not reported here due to lack of space.

Fig. 8 reports the comparison of the energy spent for each bit by an energy-aware network and a standard network. The energy per bit is computed as:²

$$EB(t) = \frac{\sum_i P_{on}^i(t) + \sum_{ij} P_{on}^{ij}(t)}{\sum_{sd} T^{sd}(t)} = \frac{P_{on}^{TOT}(t)}{T} \quad (6)$$

The figure reports also the efficiency gain computed as:

$$Gain_{EB}(t) = \frac{EB_{ST}(t) - EB_{EA}(t)}{EB_{ST}(t)} \quad (7)$$

where $EB_{ST}(t)$ and $EB_{EA}(t)$ are the energy per bit at time t for a standard and the energy-aware network respectively. The plot shows that with the energy-aware network design it is possible to reduce the cost of transporting information during the whole day, with higher gains during the night time. Indeed, during the night time, the saving is higher than 30% of the energy needed to transport a single bit.

Finally, Fig. 9 shows the total traffic flowing on the network versus the total available capacity, i.e., the average network utilization:

$$\frac{\sum_{sd} T^{sd}(t) H^{sd}(t)}{\sum_{ij} C_{on}^{ij}(t)} \quad (8)$$

²Notice that $1W = 1J/s$, so that the unit of measurement of energy per bit is J/b .




Fig. 8. Energy efficiency comparison between the energy-aware network and a standard network




Fig. 9. Total traffic versus total capacity with the energy-aware network and a standard network

where $H^{sd}(t)$ is the length of the shortest path from s to d at time t . The plot shows that the energy-aware network utilization is much higher, since the spare capacity is reduced to the minimum. Due to the link capacity granularity, however, during the night the network utilization is very low, since the link offered load is very low. This makes it worth investigating possible solutions in which link capacities can be reduced during off-peak time, in order to reduce the waste of energy.

B. Real Profile

In this section, we use the real traffic profile reported in Fig. 4 to evaluate the actual maximum power saving that can be achieved by the specific network scenario we are considering. For the sake of brevity, we report only the energy efficiency gain in Fig. 10. Results confirm that in realistic cases, the energy gain per bit is always higher than 12%, with top gain reaching 33% of saving.

TABLE II
ENERGY SUMMARY WITH THE REAL TRAFFIC PROFILE

Traditional	Energy-Aware	Saving
12.47 GWh/year	9.54 GWh/year	23.5%

To compute the total saving per year, Table II shows the energy required by a traditional network and an energy-aware network. Note that the daily energy consumption E_{day}^{TOT} has




Fig. 10. Energy efficiency gain (Real Profile)

been computed from $P_{on}^{TOT}(t)$ by:

$$E_{day}^{TOT} = \int_{t=0}^T P_{on}^{TOT}(t) dt \approx \sum_{i=0}^{N-1} P_{on}^{TOT}[t_i] \Delta t_i \quad (9)$$

Notice that since $T = 24h$ the measurement unit for power is kWh . By assuming that the traffic profile is repeated over the days, we can compute the total energy consumption in one year as $E_{year}^{TOT} = E_{day}^{TOT} \times 365$.

We can see that the power consumption of a traditional network is huge, but with an energy-aware approach it is possible to save more than 23% of energy in a year.

V. CONCLUSIONS AND FUTURE WORK

In this paper we have considered a realistic IP network topology and evaluated the amount of energy that can be potentially saved when nodes and links in the network are turned off during off-peak periods. A simple algorithm has been proposed to select the network equipments that must be powered on in order to guarantee the service. Results show that it is possible to save more than 23% of total energy consumption, which corresponds to a saving of 3GWh/year.

ACKNOWLEDGMENT

The work described in this paper was performed with the support of the BONE project (Building the Future Optical Network in Europe), funded by the European Commission through the 7th Framework Programme, and the WiFi4Energy project, funded by Regione Piemonte.

REFERENCES

- [1] Global Action Plan Report, *An inefficient truth*, <http://www.globalactionplan.org.uk/>, 2007
- [2] SMART 2020 Report, *Enabling the low carbon economy in the information age*, <http://www.theclimategroup.org>, 2008.
- [3] H. O. Scheck, J. Louhi, *Energy Efficiency of Cellular Networks*, W-GREEN 2008, Finland.
- [4] R. Tucker, J. Baliga, R. Ayre, K. Hinton, W. Sorin, *Energy Consumption in IP Networks*, ECOC Symposium on Green ICT 2008, Belgium.
- [5] S. Pileri, *Energy and Communication: engine of the human progress*, INTELEC 2007 keynote, Rome, Sep. 2007.
- [6] M. Gupta, S. Singh, *Greening of the Internet*, Proceedings of ACM SIGCOMM, Karlsruhe, Germany, August 2003.
- [7] P. Barford, J. Chabarek, C. Estan, J. Sommers, D. Tsang, S. Wright, *Power Awareness in Network Design and Routing*, IEEE INFOCOM 2008, Phoenix, USA, April 2008.
- [8] L. Chiaraviglio, M. Mellia, F. Neri, *Reducing Power Consumption in Backbone Networks*, IEEE ICC 2009, Dresden, Germany.