



# Energy-aware routing in IP networks: A reality check

Aruna Prem Bianzino<sup>1</sup>, Claude Chaudet<sup>1</sup>, Federico Larroca<sup>2</sup>, Dario Rossi<sup>1</sup>, Jean-Louis Rougier<sup>1</sup>

<sup>1</sup>Institut TELECOM, TELECOM ParisTech, CNRS LTCI UMR 5141, Paris, France

firstname.lastname@enst.fr

<sup>2</sup>Instituto de Ingeniería Eléctrica, Facultad de Ingeniería, Universidad de la República, Uruguay

flarroca@fing.edu.uy

**Abstract**— In this work, we analyze the design of green routing algorithms and evaluate the achievable energy savings that such mechanisms could allow in several realistic network scenarios. We formulate the problem as a minimum energy routing optimization, which we numerically solve considering a core-network scenario, which can be seen as a worst-case for energy saving performance (as nodes cannot be switched off). To gather full-relief results, we analyze the energy savings in various conditions (i.e., network topology and traffic matrix) and under different technology assumptions (i.e., the energy profile of the network devices).

These results give us insight into the potential benefits of different “green” technologies and their interactions. In particular, we show that depending on the topology and traffic matrices, the optimal energy savings can be modest, partly limiting the interest for green routing approaches for some scenarios. At the same time, we also show that the common belief that there is a trade off between green network optimization and performance does not necessarily hold: in the considered environment, green routing has no effect on the main network performances such as maximum link utilization.

## I. INTRODUCTION

Consciousness on energy consumption is nowadays rising in the ICT field, the network being an important contributor to the total power consumption: the effort of bringing energy-awareness in network elements and processes is usually referred to as *green networking*.

Once a network has been designed (i.e., the resources that will compose it have been deployed), a periodical off-line process is applied to optimize the utilization of resources, which we will refer to as “routing optimization”. This classical process consists in particular in determining the paths used for each origin-destination pair or, equivalently, to ingress-egress routers in a transit network. Common optimization objective is to avoid congestion by e.g., balancing the traffic as evenly as possible on the network links, or by ensuring that maximum link utilization always remains below a given threshold. In pure IP networks, the path used by each flow

is determined by the Internal Gateway Protocol (IGP), based on link administrative weights. Network dimensioning is thus handled by careful weight assignments, for instance using IGP Weight Optimization (IGP-WO) algorithms [1].

One of the most common green practices in network dimensioning consists in *resource consolidation*: this technique aims at reducing the energy consumption due to devices underutilized at a given time. Given that the traffic level in a given network approximately follows a well known daily and weekly behavior, there is an opportunity to aggregate traffic flows over a subset of the network devices and links, allowing other devices to be temporarily switched off. This solution shall of course preserve connectivity and Quality of Service (QoS), for instance by limiting the maximum utilization over any link. In other words, the required level of performance will still be guaranteed, but using an amount of resources that is dimensioned over the actual traffic demand, rather than for the peak demand. Flow aggregation may be achieved, for example, through a proper configuration of the routing weights in an IP network.

This approach has been evoked in [2] as a hypothetical working direction. The authors of [3] take a first practical step in this direction, with the proposal and evaluation of some greedy heuristics, that are based on the ranking of nodes and links with respect to the amount of traffic that they would carry in an energy-agnostic configuration.

In this work, we instead formulate the green routing as an optimization problem, which we numerically solve to evaluate the achievable energy savings. Aiming at a realistic evaluation, we consider (i) several power models corresponding to different technologies, (ii) an actual network topology and (iii) real traffic matrices, taken from an operational network. At the same time, we take a deliberately conservative approach by choosing a core-network where, since all nodes generate and receive traffic, no node can be turned off – which constitutes a realistic worst-case scenario for network resource consolidation scheme. Our goal is to get insight into potential

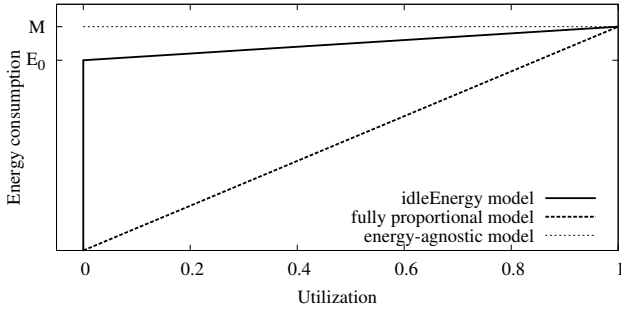


Fig. 1. The different models for the network device energy consumption, expressed as parameterized function of the device utilization.

TABLE I

ENERGY CONSUMPTION PARAMETERS IN WATTS, FOR THE DIFFERENT NETWORK ELEMENTS.

Network element	$E_0$ [Watt]	$M$ [Watt]	Ref.
Nodes	$0.85C^{3/2}$	$C^{3/2}$	[6]
(0-100) Mbps links	0.48	0.48	[7], [4]
(100-600) Mbps links	0.90	1.00	[7], [4]
(600-1000) Mbps links	1.70	2.00	[5]

energy savings in realistic scenarios, and to identify room for future improvement.

## II. ENERGY MODEL

In order to evaluate the energy saving of our green solution, it is fundamental to rely on an accurate energy consumption model. Yet, we point out that obtaining energy consumption figures for real network infrastructures represents a very challenging task (due to the inconsistency of the different models, which further become quickly out-of-date). We thus take special care in the definition of a general model, describing the devices' energy consumption as a function of their utilization. The model is expressed in a parametric form, that makes it easily extensible to other cases, and which we tune in this work according to power figures available in [4]–[7].

It is generally accepted that network device energy consumption grows linearly between a minimum value  $E_0$ , which corresponds to the idle state, and a maximum value  $M$ , which corresponds to the maximum utilization [8]. Furthermore, a null energy consumption is assumed when the device utilization is equal to 0, in which case the device is set to a sleeping state. We refer to this model as “idleEnergy”, which is illustrated in Fig. 1 by a solid line.

For what concerns the actual values of parameters  $E_0$  and  $M$ , we rely on the mostly accepted and diffused energy figures available in the literature. Table I summarizes the parameters we used, where  $C$  represents the node switching capacity. As the overall switching capability for nodes in the considered topology is not available, we considered a node as being able to switch the double of the sum of the capacity of all links connected to it. This is a design conservative choice, that would allow the network manager to add a reasonable number of links without having to change the devices.

Two special cases of this energy model are of particular interest in our analysis. In the *fully proportional* model, the

parameter  $E_0$  is equal to 0. This model represents an ideal case where energy consumption varies linearly with the device utilization, between 0 and  $M$ . This model is illustrated in Fig. 1 by a dashed line. It represents the behavior of fully *energy-aware* devices, such as communication links supporting rate adaptation [9]. Nodes could also present such a behavior when their components are regulated in function of the load (e.g., Dynamic Voltage Scaling (DVS), modular switching fabrics, etc.). The fully proportional model is thus a resultant of several green technologies, which are not necessarily available today, and is thus to be considered as a futuristic scenario. On the opposite, in the *energy agnostic* model the  $E_0$  parameter is equal to  $M$ , as illustrated in Fig. 1 by a dotted line. This case models network elements whose energy consumption is constant, independently from their load, and are never powered down (i.e., the common case today).

## III. PROBLEM FORMULATION

We represent the network as a directed graph,  $G = (N, L)$ , with  $N$  the set of nodes modeling interconnection devices and  $L$  the set of arcs modeling the communication links. For any network element  $a$  (node or link), we will denote by  $l_a$  its load and by  $c_a$  its capacity, i.e., the maximum load it can support.

Our objective is to find the network configuration (i.e., the loads and the on/off status of the nodes and links of the network) that minimizes the total network energy consumption, expressed as the sum of the consumptions of all nodes and links. Considering the model introduced in Sec. II, the consumption of each element corresponds to an affine function of the element usage, i.e., the ratio between its load and its capacity. The constant term of this affine is equal to  $E_0$  when the element is switched on and is null otherwise. To model this, we denote by the binary variable  $x_a$  the status of element  $a$  ( $x_a = 1$  whenever  $a$  is on and  $x_a = 0$  otherwise). The slope of the affine function corresponding to element  $a$  is denoted  $E_{fa}$ . Finally, links are full duplex and they are considered entirely powered as soon as one direction conveys traffic. Since in the above graph formulation the two directions are separately modeled, the link load is the sum of both directions loads. With this model, the network total energy consumption may be represented by the following expression (where the first sum needs to be divided by a factor 2 in order to avoid counting links twice):

$$\frac{1}{2} \sum_{(i,j) \in L} \left( \frac{(l_{ij} + l_{ji})E_{fij}}{c_{ij}} + x_{ij}E_{0ij} \right) + \sum_{n \in N} \left( \frac{l_n E_{fn}}{c_n} + x_n E_{0n} \right) \quad (1)$$

The load imposed to this network is defined by a traffic matrix that specifies, for every couple of ingress and egress nodes  $(s, d)$ , the traffic flowing from  $s$  to  $d$ , denoted by  $r_{sd}$  hereafter. This flow from  $s$  to  $d$  is routed across the network, generating a traffic of  $f_{ij}^{sd}$  over any link  $(i, j)$ . This traffic matrix defines the following set of constraints:

$$\sum_{(i,s,d) \in N^3} f_{ij}^{sd} - \sum_{(i,s,d) \in N^3} f_{ji}^{sd} = \begin{cases} r_{sd} & \forall (s,d) \in N^2, j = s \\ -r_{sd} & \forall (s,d) \in N^2, j = d \\ 0 & \forall (s,d) \in N^2, j \neq s, d \end{cases} \quad (2)$$

As mentioned above, to preserve QoS, no links should reach a 100% utilization, or more in general, an arbitrary value  $\alpha$  that the network operator considers safe enough. This defines the following set of constraints:

$$\sum_{(s,d) \in N^2} f_{ij}^{sd} = l_{ij} \leq \alpha c_{ij} \quad \forall (i,j) \in L \quad (3)$$

We further assume node load to be directly proportional to the traffic entering and leaving the node. In particular, we consider that they are equal, which adds the following constraints to our problem:

$$l_n = \sum_{(i,n) \in L} l_{in} + \sum_{(n,i) \in L} l_{ni} \quad \forall n \in N \quad (4)$$

Finally, we consider that a node or a link is switched off as soon as its load is equal to zero. This allows to relate variables  $x_a$  and  $l_a$  for any element of the network through the following sets of constraints:

$$Zx_{ij} \geq l_{ij} + l_{ji} \quad \forall i,j \in L \quad (5)$$

$$Zx_n \geq l_n \quad \forall n \in N \quad (6)$$

where  $Z$  is a “big” number (i.e., greater than twice the maximum between the nodes and the links capacities), used to force the variable  $x_a$  to take the value 1 when  $a$  has a load greater than 0, and the value 0 when  $l_a = 0$

Minimizing the total energy consumption (1) while satisfying all the constraints mentioned in this section is a mixed integer program, with binary variables ( $x_a$ ) and continuous variables ( $l_a$ ).

#### IV. EXPERIMENTAL RESULTS

A known problem in the evaluation of energy saving solutions is the lack of standard conditions and metrics. At the same time, a major concern in the green IT field is the ability to quantify the achievable energy reduction in a scenario that is as relevant and as objective as possible – as otherwise the promised energy gain could be as extraordinary as, unfortunately, highly unrealistic.

For this reason we decided to evaluate the potential benefits of energy-aware routing on solution using realistic, and publicly available, data. The setting we choose constitutes a worst-case scenario, so that we are able to estimate a reliable and conservative lower-bound on the achievable gain (Sec.IV-A). On such realistic scenario, we also elaborate on further properties of the solution, so to assess the impact of energy-aware routing on the achievable quality of service, possibly imposing maximum load on individual links to ensure further robustness (Sec.IV-B). Finally, as we are aware that worst-case scenario may provide a too pessimistic lower bound, we perform a careful sensitivity analysis of the solution, widening the boundaries of our investigation through careful transformation of the real input traffic matrix (Sec.IV-C).




Fig. 2. A representation of the GEANT network topology used in the solution evaluation (different link colors represent different utilization levels)




Fig. 3. Energy consumption in Watts for the different routing algorithms, applied to the different energy models.

##### A. Worst-case scenario

As input of our reality-check, we selected the GEANT topology [10], which represents a real and fairly complex network: as Fig. 2 depicts, it includes 23 nodes and 74 links. As traffic data, we selected a subset of the available Traffic Matrices (TMs), specifically 24 TMs, taken at hourly intervals between 00:30 and 23:30 of 5/5/2005. Notice that this TM set includes the complete traffic variations of a standard working day.

As a performance metric, we selected the percentage of energy saved with respect to a routing configuration using the IGP Weight-Optimization (IGP-WO) algorithm [11]. IGP-WO is the standard practice in the operator networks, we will refer to this reference scenario as “IGP-WO routing”.

We model the optimization problem with AMPL [12], and use CPLEX [13] for its numerical solution. Detailed results obtained for the three considered energy models are summarized in Tab. II, while Fig. 3 offers a graphical view of the energy saving, separately considering the node and link contributions. Both Tab. II and Fig. 3 report values averaged over the 24 TMs set.

Regarding the *idleEnergy* model, we can expect energy

TABLE II

ENERGY CONSUMPTION IN WATTS FOR DIFFERENT ROUTING ALGORITHMS, APPLIED TO DIFFERENT POWER MODELS AND SCENARIOS, AVERAGED OVER THE FULL SET OF TRAFFIC MATRICES. FOR THE CASE OF GREEN ROUTING, THE PERCENTAGE OF ENERGY SAVING WITH RESPECT TO THE CORRESPONDING IGP-WO CASE IS REPORTED IN PARENTHESES.

Scenario	IGP-WO routing			Green routing		
	Nodes	Links	Total	Nodes	Links	Total
Energy-agnostic	7676.00	59.12	7735.12	7676.00 (-0.0%)	59.12 (-0.0%)	7735.12 (-0.0%)
idleEnergy	6565.95	46.23	6612.18	6569.22 (+0.05%)	30.34 (-34.4%)	6599.56 (-0.2%)
fully proportional	307.21	10.97	318.18	286.69 (-6.7%)	5.10 (-53.5%)	291.79 (-8.3%)

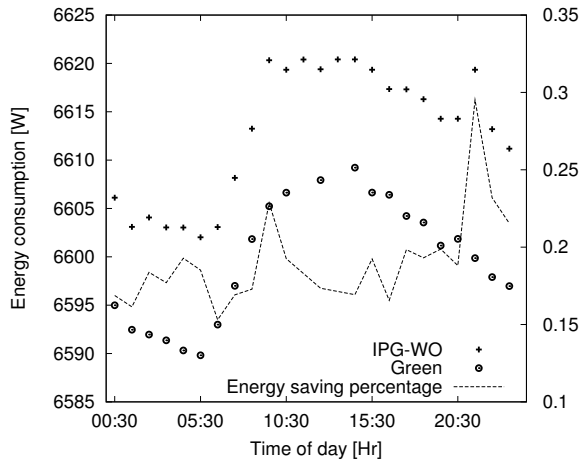


Fig. 4. Overall network energy consumption under the *idleEnergy* model, and IGP-WO and green routing cases.

savings to be mainly a consequence of switching off network elements, since this avoids the idle energy consumption  $E_0$ . Indeed, it is clear from the values considered in Tab. I that the impact of the fixed component  $E_0$  on the overall network energy consumption is much greater than the proportional energy component due to the device load ( $M - E_0$ ). Moreover, in the considered model, the energy parameters of the nodes are generally two orders of magnitude larger than the ones of the links. This means that the energy saving achievable by switching off links represents a small contribution to the total energy saving. However, given the topology and the traffic level, it is generally not possible to switch off nodes (since every node is source and destination of traffic requests), but it is possible to switch off links. The GEANT scenario represents hence a *worst-case scenario* for our solution, lower bounding the achievable energy saving.

The above considerations are verified in the results, which show a small energy saving due to nodes but a considerable one due to links, summing to a modest overall energy saving (about 0.2%). Fig. 4 shows how the network energy consumption varies for the *idleEnergy* model when considering the different traffic requests on a typical working day (i.e., over 24 hourly traffic matrices). Energy consumptions are reported for both the IGP-WO and green routing, along with the energy saving percentage.

In the case of the *fully proportional* model, the energy saving is a consequence of the aggregation of traffic over paths involving the most energy efficient devices, while we are not interested in switching off nodes and links since there is no idle

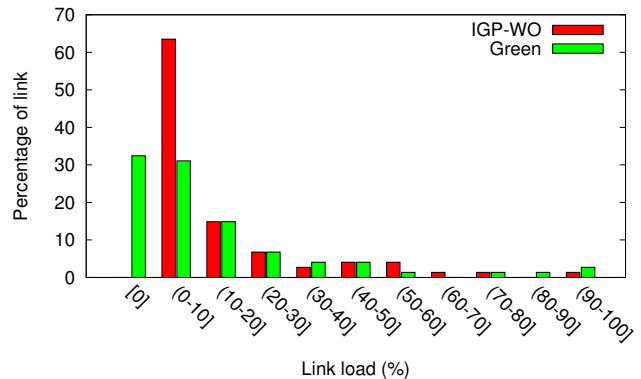


Fig. 5. Link load distribution under IGP-WO and green routing.

energy consumption ( $E_0 = 0$ ). Observing the results reported in Tab. II, we can see that it is possible to achieve a much higher energy saving by means of energy-aware devices (fully proportional model) than with nowadays devices, presenting at most a partial energy awareness (*idleEnergy* model). Therefore, we see that green routing and green technologies (such as link rate adaptation in IEEE 802.3az [14] and dynamic voltage and frequency scaling [15] techniques, which bring links and devices close to a fully proportional model) naturally interact for enhanced saving performance.

### B. QoS Considerations

In our solution, energy saving comes as consequence of switching off network elements and optimizing their utilization level with respect to their power consumption. This strategy is in opposition to the common practice to guarantee robustness and QoS in networks: i.e., redundancy of network elements and distribution of the charge over all the available paths. It is therefore imperative to analyze the variations in the network device load that energy-aware routing brings with respect to the standard IGP-WO routing case.

More precisely, we assess how a green solution displaces the load in the network to achieve energy saving, and how it affects the devices load, a performance indicator that has direct influence on users QoS. For the sake of simplicity, we report results that refer to a single scenario (namely, the 00:30 TM under the *idleEnergy* model), although qualitatively similar considerations hold for other scenarios as well. Fig. 5 reports the distribution of the link loads for *idleEnergy* model for both IGP-WO and green routing. Note that in the IGP-WO case none of the link are idle, while our solution brings a considerable number of link to a zero utilization (i.e., since




Fig. 6. Average link load for the different link types under IGP-WO and green routing.

those links are turned off to save energy). As a consequence, green routing also increases the number of links with a higher utilization level, since this is a straight consequence of aggregating the traffic on a subset of the network devices, to be able to switch off the others.

Fig. 6 shows the average link load for the IGP-WO and green routing cases: we can see that the overall average load slightly increases under green routing, and that our solution tends to move the load from the “average capacity” links, to the more energy-efficient “high capacity” links. Notice that resource consolidation is generally not possible for a number of (even less energy-efficient) “low capacity” access links, as they are located in more constrained areas of the network (i.e., at the edge, where the path diversity is lower). Notice also that the green solution does not increase considerably the average link utilization: indeed, even though resource consolidation may slightly increase the overall network load (as a consequence of longer paths), however this very same amount of traffic is generally shifted over a higher capacity link (with a thus limited impact on the link utilization).

Finally, it is interesting to further dig the solution of the optimization problem at a finer level of details, by pinpointing individual links that are switched off under green routing. This is shown in Fig. 7, where links represented as thick black lines are switched off. In the considered scenario, the switch off procedure only involves lightly loaded links: the average load on these links is 5.2% in energy-agnostic configurations. Moreover, only average and high capacity links are switched off (i.e., all link with a  $C \leq 100$  Mbps are on). Overall, as all nodes connected to a switched-off link are also connected to at least another link with at least the same capacity, we can conclude that the impact of green routing on QoS will be minimal (as the traffic will be redirected on alternative links without considerably affecting the link utilization).

Nowadays, the network operators adopt as common practice to limit the load of the links to enforce QoS and robustness in their networks. In order to reflect this practice in our solution and to obtain more realistic results, we introduce in the problem formulation (sec III) a maximum imposable load level for all the network elements, by a parameter referred to as  $\alpha$ . Fig. 8 illustrates the variation of the achievable energy saving for a range of maximum imposable load  $\alpha$ , where in this case we average results over the full set of 24 TMs




Fig. 7. Switched off network elements (in black) when green routing is applied.

(results are obtained for the fully proportional model, but considerations holds for other energy models as well). From Fig. 8 we gather that the reduction of the maximum device load does not significantly affect the achieved energy saving. This is due to the fact that the main limitation to the energy saving is represented by the topology and the traffic requests, rather than by the maximum device utilization (remember that, by design, *nodes* are never loaded more than 50%).

On the other hand, reduction of the maximum imposable load on *links* significantly affects the feasibility of the problem: it should be noticed that, in low-load scenario of 00:30 TM reported in Fig. 5, some links are loaded more than 90% under IGP-WO. Hence, reduction of the maximum device load may easily result in unfeasible solutions (the percentage of feasible solutions is reported on the right y-axis of 8).

### C. Sensitivity Analysis

Finally, we perform a sensitivity analysis of the solution by performing another set of experiments, carefully controlling the evaluation scenario.

In this case, we consider that one or more nodes of the GEANT network are no longer generating and receiving traffic, but become “core” nodes which merely route other nodes traffic.

As core nodes, we select the five most central nodes (*at1.at*, *ch1.ch*, *de1.de*, *es1.es*, and *uk1.uk*), and perform the full set of experiments by considering all combinations of  $N$  core nodes, from  $N = 1$  (i.e., 5 scenarios with a single core node) to  $N=5$  (i.e., the single scenario with all 5 core nodes). Clearly, in this case, the optimization solution may have the chance of switching off one or more of the core node, provided that routing is feasible under the reduced graph.

The aim of this set of experiments is twofold: on the one hand, we want to complement our worst-case analysis with other figures, that are at the same time realistic but controlled. On the other hand, this kind of analysis is also important as it can bring useful insights for topology design.

Results are reported in Fig. 9 for the *idleEnergy* model and the 05:30 TM, which correspond to the *minimum network load* but not to the maximum of energy saving (as can be seen



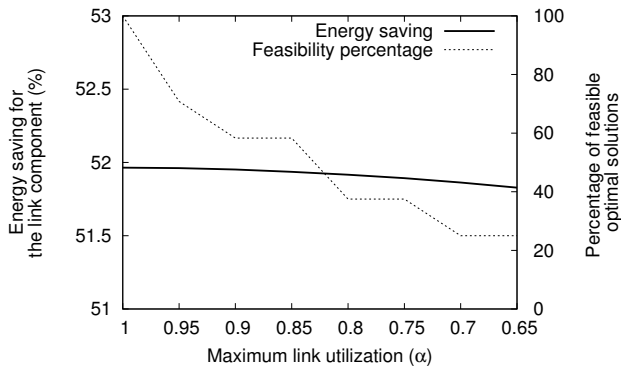


Fig. 8. Energy saving percentage against maximum imposable link utilization.

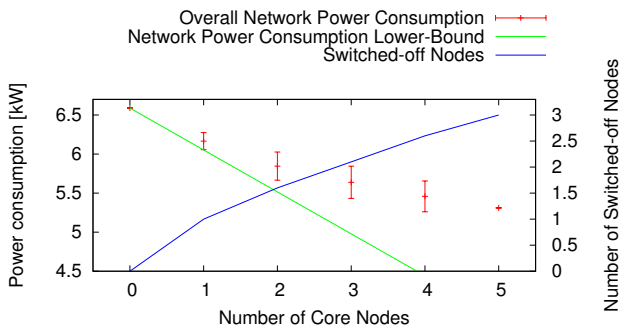


Fig. 9. Energy saving as a function of the number of core nodes.

in Fig. 4). Notice that, as the load is low, and nodes could be entirely switched off, there are opportunities for a higher energy saving: thus, the least loaded TM upper-bounds the achievable optimization gain.

Already when  $N = 1$  the total gain under the *idleEnergy* model is about 6% – corresponding to a 30-fold increase with respect to the 0.2% gain early obtained by switching off links only. Clearly, 6% roughly correspond to switching off 1/23 nodes (notice that node consumption is not homogeneous). As it can be seen in Fig. 9, this trend does not hold for growing number of core nodes, since the optimization problem is not always able to switch off all the core nodes under the routing constraints. In fact, the blue line in Fig. 9 reports on the right y-axis the average number of switched off nodes, which grows less than linearly with respect to the number of core nodes, due to the saturation of the switching capability of the network. In the same picture, the green line reports on the left y-axis a lower bound to the power consumption, obtained by switching off all the nodes which are of *core* type in the scenario.

Overall, we see that even though the network redundancy is high, energy-aware routing is effective in consolidating the unnecessary resources: thus, network topology could be dimensioned to avoid over-load for peak scenarios, and energy-aware routing could profitably be used to dynamically select the subset of necessary resources when in under-load scenario.

## V. CONCLUSION

In this work, we study energy-aware routing as an optimization problem, evaluating its effects on the network energy

consumption and on the network device load, a standard indicator for the QoS performance.

Considering the energy performance, the obtained results allow to better understand the mechanisms enabling energy saving, i.e., how the traffic is redirected to allow network devices to be switched off. Moreover, our results show that energy saving performance strongly depends on both (i) the network topology and traffic conditions and (ii) the device technology, corresponding to different power models. Indeed, green routing may provide an energy-efficient automatic adaptation of the network resources to the traffic conditions, but this should be supported by the topology design (e.g., enough path diversity) and supported by energy-efficient device design.

Considering the network QoS performance, numerical results show that, at least in the considered scenarios, achieving energy saving does not necessarily negatively affect the network performance, even if it may raise reliability issues. Instead, results also show that imposing maximum load on links may significantly limit the applicability of energy-aware routing, as many solution may become unfeasible. Thus, ISP will have to carefully select the trade off between the achievable energy efficiency gain and the robustness of the solution – as the choice of an unlucky robustness threshold may severally limit the achievable energy efficiency gains.

To the best of our knowledge, this work is the first to bring a careful and thorough reality-check on energy aware routing, by considering a publicly available real network topology, along with several traffic matrices and energy models; as a beneficial side effect, using publicly available data also promote cross comparison work. In future work, we aim at benchmarking existing heuristics such as [3] against the optimal numerical solution – both from the point of view of the achievable energy saving, and from the one of the robustness and QoS. We are also interested in further refining the energy profiles, in particular considering the underlying optical equipments. We believe that results may change significantly when taking into account optical links over the real geographical distances, as the energy consumption of long haul links would actually depend on their length (e.g., due to periodical signal regeneration).

## ACKNOWLEDGMENT

This work was funded by Eureka Celtic TIGER2 project.

## REFERENCES

- [1] B. Fortz and M. Thorup, “Internet traffic engineering by optimizing OSPF weights,” in *IEEE INFOCOM 2000*, (Tel-Aviv, Israel), 2000.
- [2] M. Gupta and S. Singh, “Greening of the Internet,” in *ACM SIGCOMM 2003*, (Karlsruhe, Germany), Aug. 2003.
- [3] L. Chiaraviglio, M. Mellia, and F. Neri, “Reducing Power Consumption in Backbone Networks,” in *IEEE ICC’09*, (Dresden, Germany), June 2009.
- [4] C. Gunaratne, K. Christensen, and B. Nordman, “Managing energy consumption costs in desktop PCs and LAN switches with proxying, split TCP connections and scaling of link speed,” *International Journal of Network Management*, vol. 15, pp. 297–310, Sept. 2005.
- [5] C. Gunaratne, K. Christensen, and S. W. Suen, “Ethernet Adaptive Link Rate (ALR): Analysis of a buffer threshold policy,” in *IEEE GLOBECOM*, (San Francisco, California, USA), Nov. 2006.
- [6] R. Tucker, J. Baliga, R. Ayre, K. Hinton, and W. Sorin, “Energy consumption in IP networks,” in *ECOC 2008*, (Brussels, Belgium), 2008.

- [7] R. Hays, A. Wertheimer, and E. Mann, "Active/Idle Toggling with Low-Power Idle." Presentation for IEEE 802.3az Task Force Group Meeting., Jan. 2008.
- [8] L. A. Barroso and U. Hölzle, "The Case for Energy-Proportional Computing," *IEEE Computer*, vol. 40, pp. 33 – 37, Dec. 2007.
- [9] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, "Reducing the Energy Consumption of Ethernet with Adaptive Link Rate (ALR)," *IEEE Transactions on Computers*, vol. 57, pp. 448–461, April 2008.
- [10] The GEANT network, "<http://www.geant.net/>."
- [11] The IGP-WO algorithm, "<http://totem.run.montefiore.ulg.ac.be/algos/igpwo.html>."
- [12] "Ampl, a modeling language for mathematical programming." <http://www.ampl.com/>.
- [13] "Ibm ilog cplex optimizer homepage." <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>.
- [14] "Ieee 802.3az task force." <http://www.ieee802.org/3/az/index.html>.
- [15] C. Isci, A. Buyuktosunoglu, C.-Y. Cher, P. Bose, and M. Martonosi, "An Analysis of Efficient Multi-Core Global Power Management Policies: Maximizing Performance for a Given Power Budget," in *Proceedings of IEEE/ACM International Symposium on Microarchitecture (MICRO 39)*, (Orlando, Florida, USA), pp. 347–358, IEEE Computer Society, Dec. 2006.