Minimizing ISP Network Energy Cost: Formulation and Solutions

Luca Chiaraviglio, *Member, IEEE*, Marco Mellia, *Senior Member, IEEE*, and Fabio Neri, *Senior Member, IEEE*

We would like to dedicate this paper to the memory of Fabio Neri, who suddenly passed away in April 2011.

***Abstract*—According to several studies, the power consumption of the Internet accounts for up to 10% of the worldwide energy consumption and is constantly increasing. The global consciousness on this problem has also grown, and several initiatives are being put into place to reduce the power consumption of the ICT sector in general. In this paper, we face the problem of minimizing power consumption for Internet service provider (ISP) networks. In particular, we propose and assess strategies to concentrate network traffic on a minimal subset of network resources. Given a telecommunication infrastructure, our aim is to turn off network nodes and links while still guaranteeing full connectivity and maximum link utilization constraints. We first derive a simple and complete formulation, which results into an NP-hard problem that can be solved only for trivial cases. We then derive more complex formulations that can scale up to middle-sized networks. Finally, we provide efficient heuristics that can be used for large networks. We test the effectiveness of our algorithms on both real and synthetic topologies, considering the daily fluctuations of Internet traffic and different classes of users. Results show that the power savings can be significant, e.g., larger than 35%.**

***Index Terms*—Green networks, link and node switchoff, network design.**

I. INTRODUCTION

THE STEADILY rising energy cost and the need to reduce the global greenhouse (such as CO_2) gas emission to protect our environment have turned energy efficiency into one of the primary technological challenges of our century [1]. In this context, information and communication technologies (ICT) are expected to play a major active role in the reduction of the worldwide energy requirements through the optimization of energy generation, transportation, and consumption. However, a number of studies estimate a power consumption related to ICT itself varying from 2% to 10% of the worldwide power consumption [2], [3]. 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 [4]. Moreover, among the main ICT sectors, 37% of the total ICT emissions are due to the telecommunication infrastructures and devices, while data centers and user terminals are responsible for the remaining part [4]. In Italy, for example, Telecom Italia is the second largest consumer of electricity after the National Railway system, consuming more than 2 TWh per year [5], which is equivalent to the energy consumed by 660 000 families in one year [6]. Similar and even more pessimistic considerations also hold for the other developed countries (see, for example, [7] for the case of Japan).

To this extent, networking devices like IP routers consume the largest majority of energy [8]. It is therefore not surprising that researchers, manufacturers, and network providers are spending significant efforts to reduce the power consumption of ICT systems from different angles.

In this paper, we aim at studying how to reduce the overall power consumption of an Internet service provider (ISP) backbone network, considering it as a single, large, and distributed system. We do not focus on reducing the power consumption of each device, but rather we aim at controlling the whole network, so as to find the minimum set of devices that must be used to meet the actual traffic demand. Traditionally, networks have always been designed to meet a given traffic demand, e.g., peak-hour traffic under quality-of-service (QoS) constraints such as maximum links load or robustness to device failures. Minimizing the equipment investment (capital expenditure, or CAPEX) has been the traditional objective function. The intuition, however, suggests that this approach, while being optimal during peak-hour traffic, results in an overprovisioning of capacity, causing waste of resources, including the energy used to keep the whole network up and running. In addition, the coarse granularity of today's transmission technology often forces the network providers to install high-bandwidth links that are lightly loaded most of the time. Spare resources are also present to provide a reliable service, so that additional links and nodes guarantee to recover from occasional failures. Keeping all these additional resources always powered on is a clear waste of energy, needlessly augmenting the operational expenditures (OPEX).

We consider the problem of ISP network energy consumption, facing a scenario in which network devices can be selectively turned on and off to meet the actual traffic demands and can be quickly activated in case of failures. We aim at answering

---

Manuscript received October 25, 2010; revised May 03, 2011; accepted June 28, 2011; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor O. Bonaventure. This work was supported by the European Community's Seventh Framework Programme under Grant Agreement 257740 (Network of Excellence "TREND"). A preliminary version of this paper appeared in the IEEE International Conference on Communications (ICC), Dresden, Germany, June 14–18, 2009, and the 1st International Workshop on Green Communications (GreenComm), Dresden, Germany, June 18, 2009.

L. Chiaraviglio and M. Mellia are with the Dipartimento di Elettronica, Politecnico di Torino, Turin 10129, Italy (e-mail: chiaraviglio@tlc.polito.it; mellia@tlc.polito.it).

F. Neri, deceased, was with the Dipartimento di Elettronica, Politecnico di Torino, Turin 10129, Italy.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2011.2161487

questions like the following: "What is the minimum set of links and routers[1] that can support current traffic demands and that minimize the total power consumption? Is it possible to save energy by increasing the maximum load each link can tolerate?

This intuition has been already proposed in the literature (see Section VII for a broader discussion), starting from the pioneering work of Gupta [9] or more recently in [10], [11], [12], and [13]. In this paper, we further push this intuition by: 1) precisely formulating the problem; 2) devising optimal solutions that can scale up to middle-sized networks; 3) providing efficient heuristics in case larger networks are considered; and 4) assessing the effectiveness of the proposed approach considering a real test case, showing that power saving can be significant, e.g., larger than 35%.

The paper is organized as follows. The mathematical formulation of the problem and the proposed algorithms are presented in Sections II and III. The description of the topologies used for performance assessment is reported in Section IV. Section V details the results obtained. An implementation discussion is reported in Section VI. Section VII overviews the related works. Finally, conclusions are drawn in Section VIII.

II. PROBLEM FORMULATION

We consider an ISP network, in which access nodes (e.g., DSLAMs, Node-B's in 3G networks, optical line terminals in passive optical networks) aggregate users' traffic. Access nodes are sources and destinations of information and are connected to the ISP transport network, whose topology is a generic mesh. Transport nodes are neither sources nor destinations of traffic. The traffic matrix at each time is assumed to be known. Considering energy consumption model of devices, we simply assume that the power consumed is independent from current load, so a constant amount of energy is consumed when a device is on. This assumption is representative of current network devices, as reported by real measurements [10].

An informal description of the design problem we consider is the following.

> Given: 1) a physical network topology comprising routers and links, in which links have a known capacity; 2) the traffic demand exchanged by all source/destination node pairs at a given time; 3) the power consumption of each link and router,
>
> Find the set of routers and links that must be powered on so that the total power consumption is minimized,
>
> Subject to flow conservation and maximum link utilization constraints.

More formally, we can provide an integer linear programming (ILP) formulation of the problem to precisely define it.

Let us represent the network infrastructure as a di-graph $G=(V,E)$, where V is the set of vertices and E is the set of edges. Vertices represent network nodes, while edges represent network links, being $N=|V|$ and $L=|E|$ the number of nodes and links, respectively. Let c_{ij} be the capacity of the link from node i to node j, and let $\alpha\in[0,1]$ be the maximum link utilization that can be tolerated. Let t^{sd} be the average amount of traffic going from node $s=1,\dots,N$ to node $d=1,\dots,N$, i.e., $\{t^{sd}\}$ represents the traffic demand.

Let $x_{ij}\in\{0,1\},\ i=1,\dots,N,\ j=1,\dots,N$ be binary variables that take the value of 1 if the link from node i to node j is present in the topology and powered on. Similarly, let $y_i\in\{0,1\},\ i=1,\dots,N$ be binary variables that take the value of 1 if node i is powered on. Let $f_{ij}^{sd}\in[0,t^{sd}]$ denote the amount of flow from s to d that is routed through the link from i to j. Similarly, let f_{ij} be the total amount of traffic flowing on the link from i to j with capacity c_{ij}. Finally, let $\mathcal{PL}_{ij}$ and $\mathcal{PN}_i$ be the power consumption of link from i to j, and of node i, respectively.

Given the previous notations, we provide different versions of the optimization problem formulation. We start from a simple and complete formulation that can be solved only for trivial cases. We then elaborate on it, deriving more complex formulations that can be solved to the optimum for medium-sized networks, i.e., up to hundreds of nodes.

A. Basic Formulation (OPT-V0)

Minimize:

$$\mathcal{P}_{\text{tot}}=\sum_{i=1}^{N}\sum_{j=1}^{N}x_{ij}\mathcal{PL}_{ij}+\sum_{i=1}^{N}y_i\mathcal{PN}_i \tag{1}$$

Subject to:

$$\sum_{j=1}^{N}f_{ij}^{sd}-\sum_{j=1}^{N}f_{ji}^{sd}=\begin{cases}t^{sd}, & \forall s,d,i=s\\ -t^{sd}, & \forall s,d,i=d\\ 0, & \forall s,d,i\neq s,d\end{cases} \tag{2}$$

$$f_{ij}=\sum_{s=1}^{N}\sum_{d=1}^{N}f_{ij}^{sd}\qquad\forall i,j \tag{3}$$

$$f_{ij}\le\alpha c_{ij}x_{ij}\qquad\forall i,j \tag{4}$$

$$\sum_{j=1}^{N}x_{ij}+\sum_{j=1}^{N}x_{ji}\le My_i\qquad\forall i. \tag{5}$$

Equation (1) minimizes the total power consumption of the network. Equation (2) states the classical flow conservation constraints, according to which traffic flows are routed using a fluid model, so that several paths can be used to transport traffic from a source until the destination node is reached. Equation (3) evaluates the total flow routed on each link. Constraint (4) forces the link load to be smaller than the maximum target utilization α, while constraint (5) states that a node can be turned off only if all incoming and outgoing links are actually turned off. The big-M method is used to force this constraint, taking $M\ge 2N$.

The presented formulation falls in the class of capacitated multicommodity minimum-cost flow problems (CMCF) [14], i.e., the problem in which multiple commodities have to be routed over a graph with capacity constraints. CMCF problems are known to be NP-hard, and therefore finding the optimal solution becomes impractical even for small networks.

Considering the complexity of this formulation, (2) entails N^4 variables, so that the problem size grows as $O(N^4)$, making it possible to solve only for trivial cases.

---

[1] In this paper, we interchangeably use the terms "node" and "router."

B. Aggregated Flows (OPT-V1)

To reduce the problem size, we can use aggregate variables to compute flow routing. More precisely, let $f_{ij}^{s} \in [0, t^{s\cdot}]$ denote the aggregate amount of flow coming from source s that is routed through the link from i to j. $t^{s\cdot} = \sum_{d=1}^{N} t^{sd}$ is the total traffic injected by source node s into the network. Flow conservation (2) and (3) can then be rewritten as

$$\sum_{j=1}^{N} f_{ij}^{s} - \sum_{j=1}^{N} f_{ji}^{s} = \begin{cases} t^{s\cdot}, & \forall s, i = s \\ -t^{si}, & \forall s, i \neq s \end{cases} \tag{6}$$

$$f_{ij} = \sum_{s=1}^{N} f_{ij}^{s} \qquad \forall i, j. \tag{7}$$

Equations (1), (4), and (5) are unchanged.

According to this formulation, the actual routing information for each (s, d) is lost, but the number of variables is reduced to $O(N^3)$ thanks to the aggregate form. Once x_{ij} and y_i have been found, routing information can be computed by reverting to the disaggregate variable f_{ij}^{sd}, and considering (2) and (3). This latter problem involves real variables only and can be easily solved [15]. Despite the reduced complexity, the aggregate formulation problem can still be solved to the optimal in trivial cases only.

C. Additional Constraints (OPT-V2)

We now introduce some additional constraints to limit the region of admissible solutions, therefore improving the convergence time of the solver. We start by explicitly stating that access nodes are traffic sources and sinks, and therefore they cannot be powered off, i.e.,

$$y_i = 1 \qquad \forall i = s, \quad \forall i = d. \tag{8}$$

Then, we explicitly exploit that, for protection purpose, access nodes are typically multihomed, i.e., they are connected to two distinct backbone nodes that collect traffic from several access nodes in the same area. For the sake of simplicity, we consider the case in which access nodes are linked to two distinct backbone nodes. This approach can be extended also to the generic multihomed case. Let us define two backbone nodes i and j to be *adjacent* if there exists an access node that is connected to both of them, i.e.,

$$\begin{aligned} &\mathrm{adj}(i,j)\{(i,j) \in E : \\ &\qquad \text{if } \exists d : x_{id} = x_{jd} = 1 \vee \exists s : x_{si} = x_{sj} = 1\}. \end{aligned} \tag{9}$$

Then, we can impose that at least one among the set of adjacent nodes must be powered on, i.e.,

$$y_i + y_j \geq 1 \qquad \forall i, j \in \mathrm{adj}(i,j) \tag{10}$$

Note that both (8) and (10) increase linearly with the number of nodes N, and the size of the problem is still heavily influenced by the flow conservation constraints (6).

D. Reduced Form (OPT-V3)

We further improve the formulation by aggregating all the source and destination nodes that are linked to the *same* backbone nodes. As in the previous formulation, we consider the case in which access nodes are dual-homed to backbone nodes. Fig. 1 sketches the reduction procedure. Let consider two backbone nodes i and j (nodes at the bottom in Fig. 1). Let $\mathcal{K}^{ij}$ be the set of access nodes that are connected to i and j only. It is possible to aggregate nodes in $\mathcal{K}^{ij}$ into a single node r that is still dual-homed to i and j using two virtual links (r, i) and (r, j). For ease of notation, we use $\mathcal{K}$ in the following.

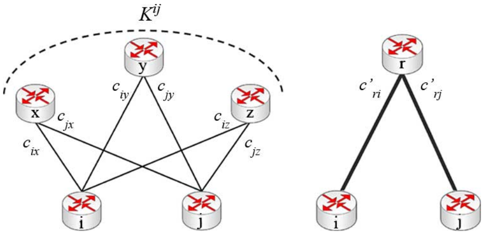

Fig. 1. Access nodes x, y, z are connected to the same two backbone nodes i, j on the left. x, y, z are aggregated into node r on the right.

Formally, the capacities c'_{ri} and c'_{rj} of links (r, i) and (r, j) can be expressed as[2]

$$\begin{aligned} c'_{ri} &= \sum_{k \in \mathcal{K}} c_{ki} - \sum_{k \in \mathcal{K}} \sum_{k' \in \mathcal{K}} t^{kk'} \qquad \forall i \in K \\ c'_{rj} &= \sum_{k \in \mathcal{K}} c_{kj} - \sum_{k \in \mathcal{K}} \sum_{k' \in \mathcal{K}} t^{kk'} \qquad \forall j \in K. \end{aligned} \tag{11}$$

Similarly, the capacities c'_{ir} and c'_{jr} of links (i, r) and (j, r) can be expressed as

$$\begin{aligned} c'_{ir} &= \sum_{k \in \mathcal{K}} c_{ik} - \sum_{k \in \mathcal{K}} \sum_{k' \in \mathcal{K}} t^{kk'} \qquad \forall i \in K \\ c'_{jr} &= \sum_{k \in \mathcal{K}} c_{jk} - \sum_{k \in \mathcal{K}} \sum_{k' \in \mathcal{K}} t^{kk'} \qquad \forall j \in K. \end{aligned} \tag{12}$$

To compute the traffic node r sends and receives, we have

$$t^{rd} = \sum_{s \in \mathcal{K}} t^{sd} \qquad \forall d \notin K \quad t^{sr} = \sum_{d \in \mathcal{K}} t^{sd} \qquad \forall s \notin K. \tag{13}$$

We set to zero the mutual traffic among the K nodes

$$t^{rd} = 0 \qquad \forall d \in K \quad t^{sr} = 0 \qquad \forall s \in K. \tag{14}$$

Finally, the power consumption is equal to the sum of the original ones

$$\mathcal{PN}_r = \sum_{k \in \mathcal{K}} \mathcal{PN}_k \tag{15}$$

and

$$\begin{aligned} \mathcal{PL}_{ir} &= \sum_{k \in \mathcal{K}} \mathcal{PL}_{ik} \quad \mathcal{PL}_{ri} = \sum_{k \in \mathcal{K}} \mathcal{PL}_{ki} \\ \mathcal{PL}_{jr} &= \sum_{k \in \mathcal{K}} \mathcal{PL}_{jk} \quad \mathcal{PL}_{rj} = \sum_{k \in \mathcal{K}} \mathcal{PL}_{kj}. \end{aligned} \tag{16}$$

The above procedure must be repeated for all pairs (i, j) of backbone nodes, and considering then all sets $\mathcal{K}$ of access nodes that are connected to i and j. Depending on the size of $\mathcal{K}$, the

---

[2] Notice that we compute the residual capacity when one out of two equivalent links is powered off.

complexity of the resulting problem is reduced, so it can then be solved using a standard ILP solver.

After the optimal solution of the reduced version has been found, it is possible to revert to the original access nodes (and links) by solving small and independent problems, one for each node r that has been created. In particular, node r has to be reverted back to the original nodes in $\mathcal{K}$, along with the original links. Only decision variables x_{kj} and x_{ik} for $k \in \mathcal{K}$ must be assigned using, e.g., OPT-V2 formulation.

Notice that the solutions of OPT-V3 are in general a subset of OPT-V2 (and OPT-V1) possible solutions. Consider an aggregated link (i, r). Only two cases are possible according to OPT-V3: either $x_{ir} = 0$ or $x_{ir} = 1$. This forces all $x_{ik} = 0$, $\forall k \in \mathcal{K}$, or $x_{ik} = 1$, $\forall k \in \mathcal{K}$. This leads OPT-V3 to give in general an upper bound to the solution of OPT-V2. In practice, as shown in Section V, the solution obtained using OPT-V3 has always been found to be equal to the optimal solution of OPT-V2.

E. Lower Bound

A lower bound to $\mathcal{P}_{\text{tot}}$ can be found by identifying the minimum-cost subset of edges and vertices of $G(V, E)$ such that there exists a path from any s to d for which $t^{sd} > 0$. In the basic formulation, this can be achieved by replacing (4) with

$$f_{ij} \leq M' x_{ij} \qquad \forall i, j \tag{17}$$

for any constant $M' \geq \sum_{s,d} t^{sd}$. In the realistic case in which any source node is allowed to exchange traffic with any destination node, this leads to a *Minimum Steiner Tree* problem in graphs, i.e., the problem of finding the minimum-cost undirected tree in G that spans among all source and destination nodes. The Minimum Steiner Tree problem is also an NP-hard problem, but several heuristics and approximation algorithms are available [16].

III. HEURISTIC APPROACH

Given the NP-hard formulations presented, finding the optimal solution using an ILP solver is viable up to a given number of nodes. Therefore, heuristic approaches have to be adopted if the problem size is too large. In this section, we present a set of heuristics explicitly designed to find an admissible solution.

All heuristics start by considering a network in which all elements are powered on, hence $x_{ij} = 1$ for every existing link in the considered topology, and $y_i = 1 \; \forall i$. Then, the algorithm checks iteratively if a given element (either a node or a link) can be turned off.

At each iteration, the considered element is removed from the graph, and traffic is then rerouted on the residual graph.[3] After rerouting, if (2) and the utilization constraint (4) are still fulfilled, then the selected element is definitively powered off. Algorithm 1 and Fig. 2 report a schematic description of the heuristics.

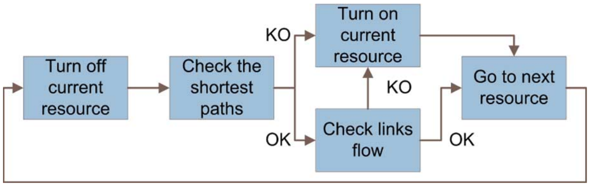

Fig. 2. Turning off technique.

Algorithm 1: Pseudocode Description of the Proposed Heuristics

```
{node optimization}
sort_nodes(node_array, order_type);
for $i = 1; i <= N; i++$ do
    disable_node(node_array ([$i$]);
    {for each $(s, d)$ pair, compute the shortest path.
    In case of tie, pick one shortest path at random}
    paths=compute_all_shortest_path();
    compute_all_link_flow(paths);
    if (check_paths(paths) == false) ||
    (check_flows(paths) == false)) then
        enable_node(node_array[$i$]);
    end if
end for
{link optimization}
sort_links(link_array, order_type);
for $(j = 1; j <= L; j++)$ do
    disable_link(link_array[$j$]);
    paths=compute_all_shortest_path();
    compute_all_link_flow(paths);
    if (check_paths(paths) == false) ||
    (check_flows(paths) == false) then
        enable_link(link_array[$j$]);
    end if
end for
```

The algorithms presented in this paper share the same intuition: The energy saving achieved by turning off nodes is higher than by switching off single links [18], and switching off a node is more difficult than switching off a single link. This suggests that the algorithm should try to turn off nodes first and then links.

The node set is first sorted considering a given rule before iterating through all the nodes. We consider the following sorting rules:

- random (R);
- least-link (LL);
- least-flow (LF);
- most-power (MP).

The random heuristic sorts nodes in random order.

The least-link heuristic sorts nodes according to the number of links that are sourced/sinked at each node, so nodes with a small number of links are considered first, i.e., V is sorted in increasing value of

$$X_i = \sum_{j=1}^{N} x_{ij} + \sum_{j=1}^{N} x_{ji}.$$

---

[3] Differently from the ILP formulation, we use a simple shortest-path algorithm to route the traffic.

The least-flow heuristic takes instead into account first the nodes with the smallest amount of information flowing through them, i.e., V is sorted in increasing value of

$$F_i = \sum_{j=1}^{N} f_{ij} + \sum_{j=1}^{N} f_{ji}.$$

Finally, the most-power heuristic tries to switch off first the nodes with the highest power consumption, i.e., V is sorted in decreasing value of $\mathcal{PN}_i$.

Considering the algorithms to turn off links, three sorting criteria are considered:

- random (R);
- least-flow (LF);
- most-power (MP);

which leverage on the same intuition of the corresponding node sorting heuristics: The random policy sorts links in random order; the least-flow policy sorts links in increasing order of carried flow, i.e., E is sorted in increasing value of f_{ij}; the most-power heuristic sorts links according to their power consumption $\mathcal{PL}_{ij}$.

All possible node/link sorting combinations have been studied. Besides these heuristics, we also tested the corresponding ones in which a decreasing order is adopted (increasing for MP). Since they all perform consistently worse (as expected),[4] we decided not to include them in this paper.

IV. TOPOLOGY DESCRIPTION

Performance of the proposed approaches is tested considering two scenarios: a test case, which considers a simplified version of the actual topology of a national ISP, and a second benchmarking data set, based on random topologies to assess the algorithm performance versus different parameters.

A. *Real Topology*

The real topology we consider has been derived from the actual topology used by a large ISP in Italy. It follows a typical hierarchical design, as reported in Fig. 3, in which four levels of nodes are present: core, backbone, metro, and access nodes.

1) Topology Description: The inner level is composed by "core nodes" [Fig. 4(a)] that are interconnected by 50-Gb/s 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 to other core nodes by two links for redundancy. Links between central POPs in different cities are about 600 km long. To offer connectivity to the Internet, a peering router is connected to four core routers by means of 100-Gb/s links.

At the second level, so-called "backbone nodes" [Fig. 4(b)] are connected to the core by 20-Gb/s links. Each backbone node is connected to two central POPs. Backbone nodes are located in "chief POPs" spread in large cities. The link length between the backbone and the core routers ranges between 50–500 km.

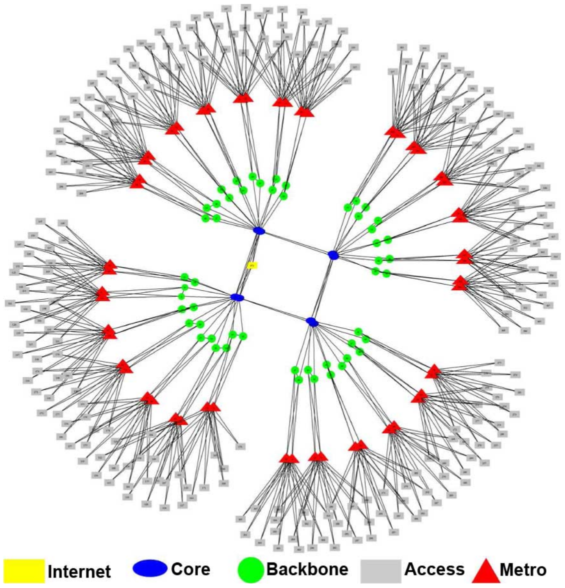

Fig. 3. ISP topology representation.

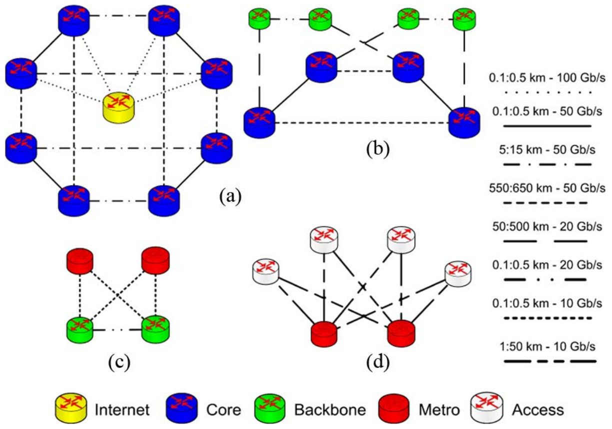

Fig. 4. Link description: (a) core, (b) backbone, (c) metro, and (d) access layers.

At the third level, "metro nodes" [Fig. 4(c)] are present. Each metro node is connected to two backbone nodes by 10-Gb/s capacity links. Metro and backbone nodes are located in the same chief POP, and links between them are then short.[5]

At the last level, there are the "access nodes" [Fig. 4(d)] that offer connectivity to the DSLAMs to which end-users are connected. Access nodes aggregate traffic from DSLAMs in the same neighborhood or small town. Each access node is dual-homed to the closest pair of metro nodes using 10-Gb/s capacity links. The length of links between access and metro ranges from 1 to 50 km.

The whole network is composed by $N = 372$ routers: 8 core nodes, 52 backbone nodes, 52 metro nodes, and 260 access nodes. Links have a cardinality equal to $L = 1436$.

2) Link Capacities: Given the coarse set of actual physical channel capacities, a link is typically formed by aggregating

---

[4] As an example, the Most Flow heuristic starts powering off first the most loaded devices, i.e., the network bottlenecks. Thus, traffic is rerouted through longer and typically uncongested paths, impairing the possibility of turning off other devices.

[5] Notice that chief POPs are composed also by other elements, e.g., the Network Access Servers (NASs) for user authentication. These devices are not considered in this paper.

several channels to form a "trunk." In particular, channel capacity is dictated according to the granularity of optical transmission systems, so that 1-, 2.5-s, 10-, and 40-Gb/s optical channels are available. $\bar{c}_{ij}$ is the capacity of a single channel, and $\lceil c_{ij}/\bar{c}_{ij} \rceil$ is the number of base-rate channels needed to actually form a trunk of capacity c_{ij}. Granularities define link routing weights as well, so the routing cost is inversely proportional to the link capacity. This is a commonly adopted policy to force the traffic to be routed through the metro and the core nodes rather than through access nodes (which are connected by means of lower capacity links).

3) Traffic Matrix: The access and the Internet Peering nodes are the only possible sources and destinations of traffic. According to real traffic estimates of the considered ISP, about 70% of the total traffic is exchanged between the Internet at large and the ISP users, while the remaining part is exchanged uniformly among the access nodes, i.e., 30% of traffic is confined within the same ISP, while 70% of traffic is coming from and going to other ISPs. Let node 0 be the peering node. Then

$$\begin{aligned} t^{s0} &= 0.7t^{s\cdot} && \forall s \neq 0 \\ t^{0d} &= 0.7t^{\cdot d} && \forall d \neq 0 \\ t^{sd} &= 0.3t^{s\cdot}/N && \forall s \neq 0, d \neq 0 \\ t^{sd} &= 0.3t^{\cdot d}/N && \forall s \neq 0, d \neq 0 \\ t^{ss} &= 0 && \forall s. \end{aligned}$$

For simplicity, for $s > 0,\ d > 0$, we assume that t^{sd} are i.i.d. random variables uniformly distributed between $[0.5, 1.5]$ units of traffic. t^{s0} and t^{0d} are uniformly distributed between $[1.16, 3.5]$. Unless otherwise specified, we consider that each link utilization cannot grow above 50% of the link capacity, i.e., $\alpha = 0.5$.[6]

The algorithm used to generate one traffic matrix operates as follows. Starting from an initial random traffic matrix $\{\hat{t}^{sd}\}$, traffic is routed through the network according to a minimum-cost path routing.[7] We then look for the mostly loaded link

$$(ij)^* = \operatorname{argmax}\left(\max_{(ij)} \frac{f_{ij}}{c_{ij}}\right)$$

from which a scaling factor

$$a = \alpha \frac{c_{ij}}{f_{ij}},\ (ij) = (ij)^*$$

is derived. Finally

$$t^{sd} = a\hat{t}^{sd} \qquad \forall s, d.$$

This means that the randomly generated traffic matrix is scaled in a way that constraint (4) holds true for all links, and that there is at least one link whose offered load is equal to the maximum one.

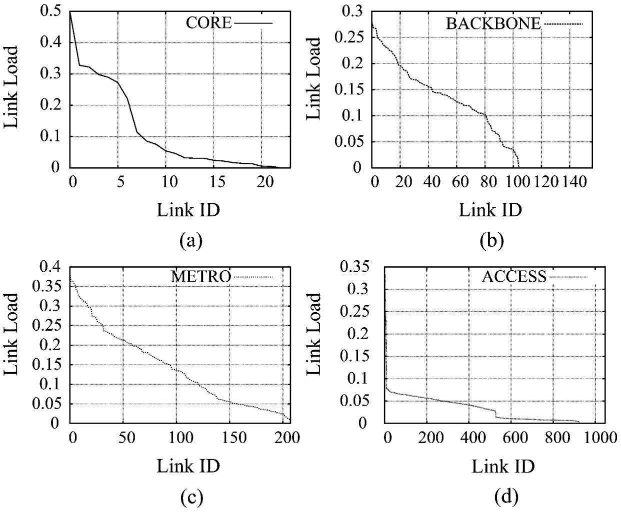

(a) (b)

(c) (d)

Fig. 5. Link load on the ISP topology for a given traffic matrix: (a) core, (b) backbone, (c) metro, and (d) access links.

We define the total traffic flowing on the network versus the total available capacity, i.e., the average network load as

$$\rho = \frac{\sum\limits_{s,d} t^{sd} H^{sd}}{\sum\limits_{i,j} c_{ij} x_{ij}} \tag{18}$$

where H^{sd} is the length of the minimum-cost path from s to d. It is immediate to see that ρ is also equal to

$$\rho = \frac{\sum\limits_{i,j} f_{ij}}{\sum\limits_{i,j} c_{ij} x_{ij}}. \tag{19}$$

The network load is an important parameter that impacts the efficiency of any green network design algorithm. Indeed, the intuition suggests that the lower is the network load, the higher is the ability to reroute traffic and to turn off network elements. The network load for the real topology scenario is about $\rho = 0.08$, which suggests that large savings are possible. To better understand the resulting link load distribution, Fig. 5 reports f_{ij}/c_{ij} when all links of the topology are present. Links are sorted by decreasing load, and grouped according to the type of node to which they are connected.[8] Link loads show a large variability, and several links are lightly loaded, resulting in a low network traffic load (see also the bottom plot in Fig. 13). This is typical in real topologies due to a general tendency to overprovisioning and to the deployment of protection resources. Access links are particularly lightly loaded, i.e., less than 0.1, while metro, backbone, and core links are more utilized on average. The bottleneck of the network is one of the peering links (for which by construction we have a link utilization equal to 50%). Still, some of the backbone and core links are lightly loaded since those are mainly protection links.

---

[6] A discussion of possible values for α is provided in Section V-B-2.

[7] In case of a tie, a random path is selected among the minimum-cost paths both to exploit network redundancy and to balance the traffic among equal-cost paths.

[8] The following hierarchy is adopted: Access links connect access nodes to metro nodes, metro links connect metro nodes to backbone nodes, backbone links connect backbone nodes to backbone nodes and core nodes, core links connect core nodes to core nodes and Internet nodes.

TABLE I
REAL TOPOLOGY: POWER CONSUMPTION

Device Type	Power Consumption	Fraction of Power
Core Nodes	10 kW	5.48%
Backbone Nodes	3 kW	11.02%
Metro Nodes	1 kW	3.66%
Access Nodes	2 kW	37.9%
Links[average]	0.6 kW	41.94%

4) Power Requirements: To model the energy consumption of routers and links, we consider the power requirements of actual devices.[9] Table I reports the power consumption considered for the different classes of nodes and the corresponding fraction of power over the total network power consumption. Notice that we ignore air conditioning costs, which can almost double 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 that takes into account possible (optical) regenerators along the optical channels, which we assume to be required every 70 km. The power consumption of the link from i to j is given by

$$\mathcal{PL}_{ij} = \left(N^{\mathrm{a}}_{ij}\mathcal{PL}^{\mathrm{a}} + \mathcal{PL}^{\mathrm{s}}\right)\lceil c_{ij}/\bar{c}_{ij}\rceil \tag{20}$$

where $N^{\mathrm{a}}_{ij} = \lfloor L_{ij}/70 \rfloor$ is the number of regenerators needed to regenerate the signal given the link length L_{ij}; $\mathcal{PL}^{\mathrm{a}}$ is the power consumption of a single amplifier, $\mathcal{PL}^{\mathrm{s}}$ is the power consumption of a single line card.

We assume that $\mathcal{PL}^{\mathrm{a}}$ is equal to 1 kW, and that $\mathcal{PL}^{\mathrm{s}}$ is equal to 100 W, based on ISP estimations. The link length is assigned as previously described and summarized in Fig. 4.

The total power consumed by the whole network amounts to $\mathcal{P}_{\mathrm{tot}} = 1.4$ MW.

B. Synthetic Topologies

To thoroughly assess the algorithm performance, we consider synthetic topologies inspired by the previously described real ISP network structure.

1) Topology Definition: We consider an ISP network, composed by nodes and optical links. All links are bidirectional, so that if link (ij) exists, then link (ji) exists as well. Three levels of nodes are considered, so backbone and core nodes in the previously described ISP topology form a single class (which we simply refer to as "core nodes"). Core nodes are located in large cities, and they are highly interconnected by means of long-distance, high-capacity links. Metro nodes are instead used to interconnect access nodes to the core nodes. Links between metro and core nodes have middle-range capacity, i.e., smaller capacity than the links interconnecting core nodes. Each metro node is connected to some of the closest core nodes and to other metro nodes. One or more metro nodes can be present in cities, and they collect traffic from access nodes spread within the city boundaries.

Access nodes, to which users are connected by means of a DSLAM or an optical line termination (OLT), are dual-homed, i.e., each access node is connected to the closest pair of metro nodes to guarantee the presence of an alternate path in case of failure. Lower-capacity links connect access nodes to metro nodes.

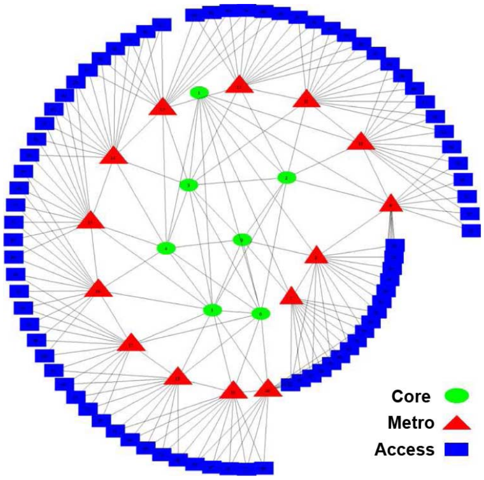

Fig. 6. Example of random topology.

Nodes are assumed to be placed on a plane. Core nodes are randomly connected to other core nodes with probability $p = 0.5$. Each metro node is then connected to the two closest core nodes and to the two closest metro nodes. Finally, access nodes are connected to the two closest metro nodes. An example of resulting topology is presented in Fig. 6. Access, metro, and core nodes are represented by squares, triangles, and circles, respectively.

Let N_{c}, N_{m}, and N_{a} be the number of core, metro, and access nodes, respectively. In the following, we identify a class of topologies by the triple $N_{\mathrm{c}}/N_{\mathrm{m}}/N_{\mathrm{a}}$.

2) Traffic Matrix and Link Capacity Assignment: Only access nodes are traffic sources and sinks. For the sake of simplicity, we consider a uniform traffic pattern, so that $t^{sd} = U[0.5, 1.5]$ Gb/s of traffic if s and d are access nodes; $t^{sd} = 0$ otherwise.

Links capacities are selected according to the following approach. As in the ISP case, three classes of links are defined: high-, middle-range-, and low-capacity links, corresponding to link interconnecting core nodes, metro nodes to core nodes and other metro nodes, and access nodes to metro nodes. Each link class has a given capacity granularity $\bar{c}_{ij}$ of 10, 2.5, and 1 Gb/s, respectively.

While in the real ISP case the traffic matrix was scaled to match the actual link capacities in the network, in this second scenario, in which both link capacities and traffic matrix are synthetic, we preferred to define a traffic matrix, and then compute link capacities to match the traffic demand. This permits to better control the load on network links, avoiding to randomly assign link capacities that can result in biased cases, e.g., access nodes connected to metro nodes by high-capacity links, but metro nodes connected to core nodes by lower-capacity links. Obviously, the load for the two network scenarios is differently distributed, with a larger average load for synthetic topologies, and this must be taken into proper account when results are analyzed.

The following procedure is used to assign link capacities. Given the traffic matrix and the minimum-cost shortest-path

---

[9] Power figures are averaged from measurements performed by the ISP.

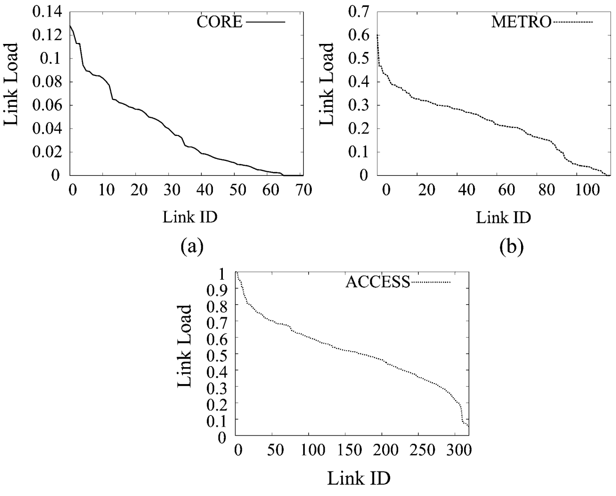

(a) (b)

(c)

Fig. 7. Example of link load on a 10/20/80 random topology: (a) core, (b) metro, and (c) access links.

TABLE II
SYNTHETIC TOPOLOGY: LINK DISTANCES

	C-C	C-M	M-M	M-A
L^{min} [km]	15	50	20	1
L^{MAX} [km]	600	500	50	50

routing, the total flows f_{ij} on network links are computed first. Then, a capacity is assigned to each link such that f_{ij} is forced to be smaller than the overprovisioning factor α, i.e.,

$$c_{ij} = \max\left(\left\lceil \frac{f_{ij}}{\alpha \bar{c}_{ij}} \right\rceil, 1\right) \bar{c}_{ij} \tag{21}$$

$\bar{c}_{ij}$ is selected according to link i, j class.

Factor α was taken equal to 1 for synthetic topologies. See Section V-B-2 for a discussion on the values of α in the ISP network. This choice of α helps increasing the average network load.

Fig. 7 shows the link load for a 10/20/80 random topology. Note again the large variability in link loads. Differently from the real ISP network, in this scenario core links are lightly utilized, while access links are more loaded. This is due to the larger connectivity among metro nodes and to the absence of Internet peering in the core (access nodes are the only possible sources and destinations for traffic in the uniform traffic matrix assumed in this scenario). Moreover, the large granularity of the core link capacity often leads to large overprovisioning.

3) Power Consumption: Power consumption for nodes is supposed to be 10, 1, and 2 kW for core, metro, and access nodes, respectively, in accordance with Table I.

The power consumption of links is modeled as in (20), where L_{ij} is assigned according to the type of node i and j, using a uniform probability between $L^{\min}$ and L^{MAX}, as reported in Table II. For example, core and metro (C-M) nodes are connected by links whose length varies from 50 to 500 km.

V. PERFORMANCE COMPARISON

We start by comparing the proposed algorithms using synthetic topologies, while leaving the real topology as a final and realistic benchmarking case. Cplex version 10.1 was used to solve the ILP formulations with a branch-and-bound algorithm, while a custom simulator was written in C to implement the heuristics. A high-end server running Linux and equipped with two AMD Athlon 64 Dual-Core processors 4200+ and 4 GB of RAM was used to obtain all results.

A. Synthetic Topologies

We start by assessing the amount of power saving that the different algorithms obtain for a given topology and traffic matrix. We also compare the computation cost of each approach to see the limits of the different ILP formulations.

We generate several topologies, considering $N_{\mathrm{c}} \in [10, 20], N_{\mathrm{m}} \in [10, 40]$, and $N_{\mathrm{a}} \in [10, 80]$. Given a topology and a traffic matrix, both the optimal formulations and the heuristics are run to evaluate the percentage of power savings that can be obtained, computed as

$$\text{Saving} = 100 \left(1 - \frac{\sum\limits_{ij} x_{ij} \mathcal{PL}_{ij} + \sum\limits_{i} y_i \mathcal{PN}_i}{\mathcal{P}_{\mathrm{tot}}}\right). \tag{22}$$

For sake of simplicity, we report only the most significant results, obtained by the original (OPT-V1) and the final reduced form (OPT-V3) formulations, as well the LF-LF and the MP-MP heuristics.

The top plot of Fig. 8 reports the total power savings. Lines are used to highlight results obtained from topologies where N_{c} and N_{m} are constant while N_{a} varies. Four $N_{\mathrm{c}}/N_{\mathrm{m}}$ pairs are considered. Power savings in the considered scenarios are significant since the network is lightly loaded. For example, for the 10/20/80 network, Fig. 7 shows that the little traffic flowing on most metro and core links could be rerouted on other links, allowing to turn off a large portion of links and nodes. When the number of access nodes increases, the power saving reduces. This is due to the larger number of sources, which increase the network load (for the load dependency on N_{a}; see the comments to the bottom plot of Fig. 10). This in turn raises the utilization of core and metro links. Therefore, it is harder to turn off a link or a node.

Considering the different algorithms, the OPT-V3 formulation saves practically the same amount of energy as OPT-V1, even if in some cases the number of access and metro links that are turned off is smaller. This is due to the final step of the OPT-V3 algorithm during which the aggregated links have to be mapped back to each original link. Fortunately, those are short-reach links, so their power consumption is almost negligible. Both LF-LF and MP-MP heuristics show very good performance for $N_{\mathrm{c}} = 10$ and $N_{\mathrm{m}} = 10$. For larger values, the LF-LF heuristic performs consistently better, with a maximum difference in power saving of 18% from the optimal solution considering the 20/40/40 scenario.

The bottom plot of Fig. 8 shows the computation times required to obtain the solution (note the logarithmic vertical scale). As expected, for the original ILP formulation, the computation time rapidly increases, becoming infeasible for large topologies. The reduced ILP form shows instead computation times below 1000 s for most topologies, thanks to the reduction

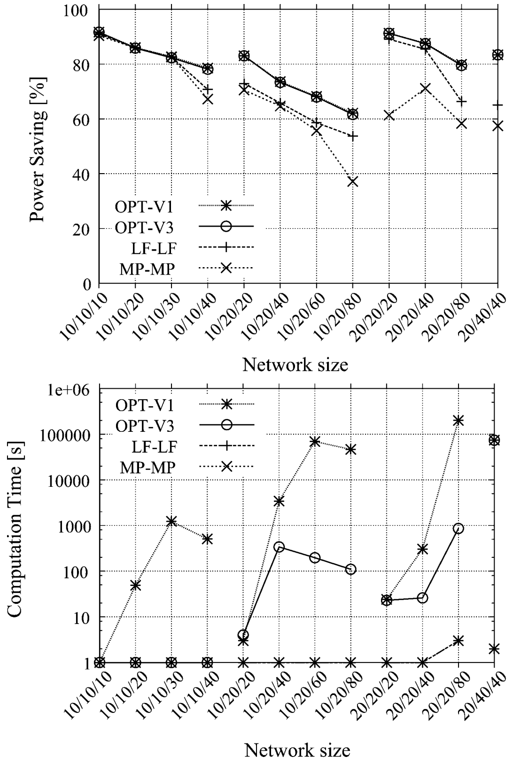

Fig. 8. *(top)* Power saving and *(bottom)* computation times considering different synthetic topologies.

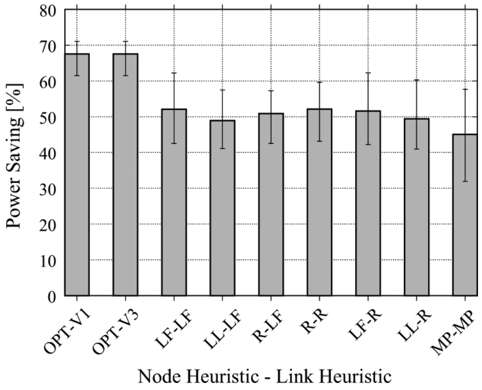

Fig. 9. Power saving considering different algorithms. 10/20/60 synthetic topologies are considered.

operated on the number of variables and constraints, which is bounded by N_{m}. Notice that both OPT-V1 and OPT-V3 require the same amount of time when aggregation is not performed, i.e., when $N_{\mathrm{m}} = N_{\mathrm{a}}$. The number of variables can be critical also in this case: For example, considering the 20/40/40 topology, the computation time is larger than 10^5 s for both OPT-V1 and OPT-V3. On the contrary, all the heuristics require less than 5 s to obtain the results in all cases.

Fig. 9 details the percentage of power saving obtained from topologies with 10/20/60 nodes. Bars report mean values, while the error bars show the minimum and maximum power saving computed over 10 independent runs. The maximum error of power saving is below 15% with 95% confidence for all the algorithms. In this case, we report results considering all the heuristics. OPT-V3 is able to guarantee the largest amount of power saving (as expected), while the other heuristics perform consistently worse, with a difference in power saving of nearly 20%. All heuristics are able to save about the same amount of power, except the MP-MP algorithm, which apparently consumes 5% of additional power on average (5% is within the confidence interval). Notice also that the MP-MP heuristic presents the largest spread between minimum and maximum saving. This is due to the greedy nature of the MP choice, which tries to first turn off the most power hungry nodes and links, possibly leading to a larger number of powered on elements.

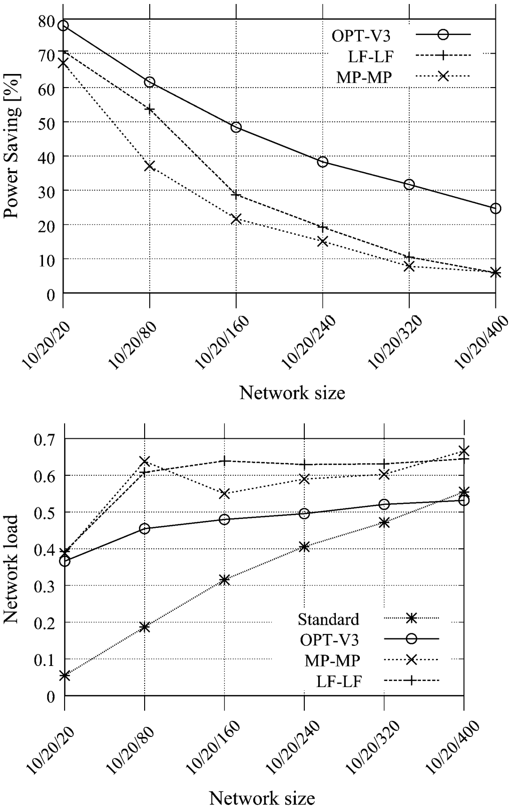

Fig. 10. *(top)* Power saving and *(bottom)* network load considering different numbers of access nodes.

To better evaluate the performance of the heuristics and the optimal algorithm, we generate networks with constant N_{c} and N_{m} but different values of N_{a}. In particular, we impose $N_{\mathrm{c}} = 10$, $N_{\mathrm{m}} = 20$, and $N_{\mathrm{a}} \in [20, 400]$.

The top plot of Fig. 10 reports the power saving. As expected, OPT-V3 takes the lead, guaranteeing a power saving between 78% and 25%. Heuristics perform quite worse, with a maximum difference in power saving of 21% for $N_{\mathrm{a}} = 320$ considering the LF-LF algorithm. Notice also that for large values of N_{a}, both the LF-LF and MP-MP heuristics save similar amounts of power. As expected, the power saving is decreasing as N_{a} increases, i.e., for increased network load.

To highlight the impact of network load on the power saving, the bottom plot of Fig. 10 reports ρ for the different scenarios. Consider first the original (labeled "Standard") network. The network load increases as N_{a} increases due to the larger number of access nodes that collect traffic from users: ρ is approximately directly proportional to N_{a}. Considering the green network solutions, turning off nodes and links increases the network load,

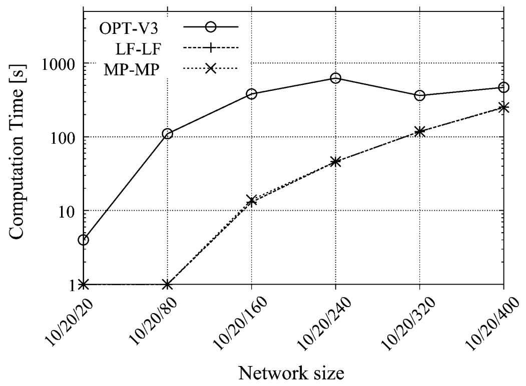

Fig. 11. Computation times considering different numbers of access nodes.

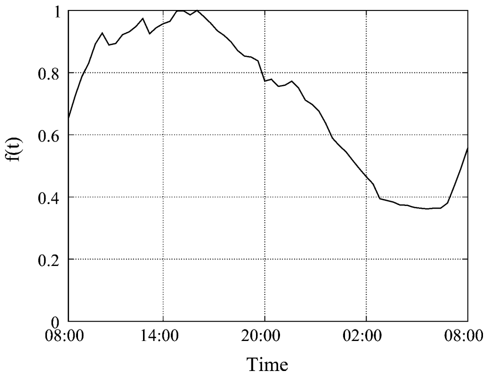

Fig. 12. Actual profile of traffic as seen on the ISP network. The same profile is assumed to affect all traffic demands.

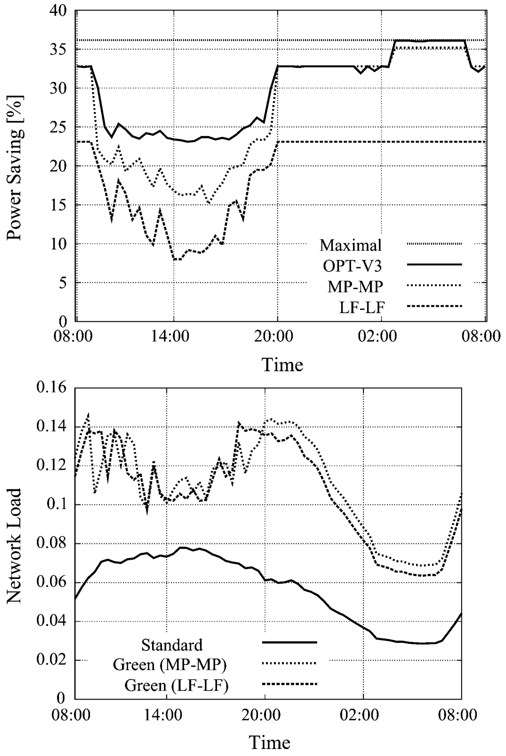

Fig. 13. *(top)* Power saving and *(bottom)* network load considering the real traffic profile.

as expected. Consequently, power savings (and correspondingly the load increase) are larger for smaller ρ since the amount of spare capacity is larger. Interestingly, the OPT-V3 network load is smaller than the heuristic one. Both LF-LF and MP-MP indeed greedily turn off backbone nodes and links (which are the lightest loaded/most power consuming resources). This causes an increase of the path length and of network load.

Finally, Fig. 11 reports the computation times for the same scenarios. OPT-V3 algorithm manages to obtain a solution in less than 500 s, thanks to the aggregate formulation in which N_a impacts only the final phase. Both LF-LF and MP-MP heuristics take consistently less time to generate a solution, but in this case the trend is increasing due to the increasing costs in computing the shortest-path algorithm each time a node or link is tested.

B. Real Case Study

We are now interested in evaluating the possible power saving that can be obtained in a realistic case. We consider the real ISP topology and traffic matrix in a scenario in which traffic varies according to a day–night pattern. Real traffic measurements collected from the ISP network are used as input data.

1) Impact of Daily Traffic Variation: We assume that the same time variation affects each traffic demand, so it can be expressed as

$$t^{sd}(t) = f(t)t^{sd} \tag{23}$$

with $f(t)$ being the shaping function at time t, and t^{sd} the traffic exchanged between s and d during peak hours. The capacity of links was assigned considering the peak-hour traffic demand. According to measurements, the shaping function (reported in Fig. 12) has a maximum equal to 1 during early afternoon and a minimum equal to 0.38 at night, i.e., off-peak traffic is 38% of the peak-hour demand.

The top plot of Fig. 13 reports the percentage of power saving using OPT-V3, MP-MP, and LF-LF. The precision of the optimal solution has been relaxed to 2% to guarantee reasonable computation times. The curve labeled "Maximal" refers to the power saving when the topology is pruned to a Steiner tree, evaluated as described in Section II-E. Interestingly, $\mathrm{Saving}(t)$ is practically constant during the night since the traffic is so little that only the connectivity constraint is effective: At night, the topology can become a Steiner tree, i.e., the achieved power saving is the absolute maximum. Saving of 36% of power is possible using the OPT-V3 algorithm. MP-MP performs well, with a maximum difference in power saving of 1% during the night, while LF-LF performs consistently worse, reaching a maximum saving of only 23%.

During the day, the power saving decreases as the traffic increases since more capacity is required to guarantee the QoS constraint. In this case, the maximum gap between OPT-V3 and the MP-MP is equal to 8%, though it is limited to the peak hours. Also in this case, the LF-LF heuristic saves less power than the other algorithms, but a minimum 8% of power saving is always possible. The better performance of the MP-MP heuristic compared to the LF-LF is due to the particular topology the ISP is using, in which the amount of overprovisioning in the

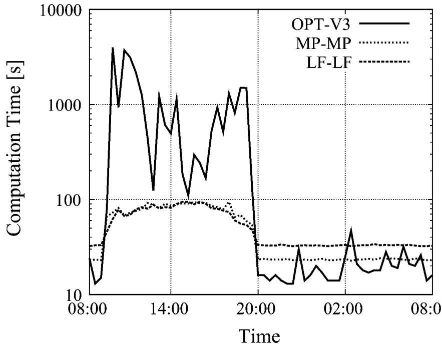

Fig. 14. Computation time considering the real traffic profile.

access and metro part of the network is very large.[10] The LF heuristic tries to first turn off all access/metro nodes and links, while moving traffic to the more power-hungry backbone nodes, leading to little power savings. On the contrary, the MP ordering criteria preferentially tries to power down core/backbone nodes and links, achieving a higher total saving.

The bottom plot of Fig. 13 shows the average network utilization ρ. As expected, more power saving is possible when the network is lightly loaded, and the green network utilization is consistently higher than for the standard network since the spare capacity is reduced. Due to unbalanced traffic, and to link capacity granularities, however, during the night the network utilization is still low (compare also with the load values in Fig. 10). This makes it worth investigating possible solutions in which link capacities can be reduced during off-peak time.

Considering the computation times, Fig. 14 shows that the OPT-V3 algorithm requires several hours to find the solution during the day. Indeed, finding good solutions is harder when the amount of overprovisioned resources is small, as confirmed also by the heuristics. Notice that without relaxing the precision to 2%, the computation times of OPT-V3 would increase to more than 10 h. During off-peak time, the set of constraints is easily met, and the optimal solution can be obtained in less than 20 s. Considering the heuristics, they both require less than 100 s, with higher times during peak hours: This is due to fact that the cost of running the shortest-path algorithm increases as the traffic increases since fewer devices are powered off.

To give more insight, Fig. 15 reports the breakdown of energy savings detailing core nodes, backbone nodes, metro nodes, and links. Values have been averaged over three different traffic matrices using the MP-MP heuristics, which proved to be the most effective one. The largest amount of savings is due to the powering off of links, with higher savings during off-peak hours. Indeed, since path redundancy is obtained by connecting each access and metro node using at least two links (dual-homing policy), it is likely that the capacity offered by just one link is sufficient to carry all traffic during off-peak hours. Also considering nodes, it is possible to turn off both core, backbone, and metro nodes during night, when the additional resources that are required to provide fault recovery do not carry traffic and can be powered off to save energy. During peak hours, the energy saving is lower, and only backbone nodes can be turned off.

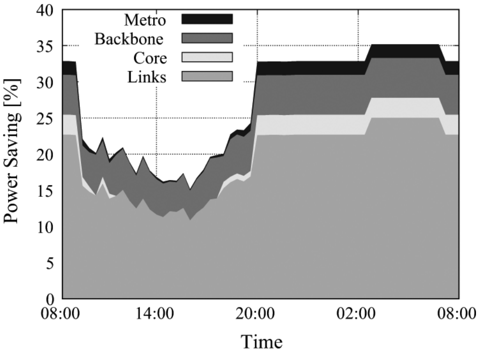

Fig. 15. Power saving breakdown considering the real traffic profile.

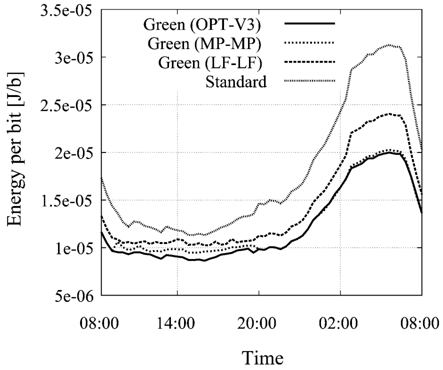

Fig. 16. Energy efficiency considering the real traffic profile.

TABLE III
ENERGY SAVINGS WITH THE REAL TRAFFIC PROFILE

OPT-V3	MP-MP
3.71 GWh/year (29.8%)	3.42 GWh/year (27.5%)

Fig. 16 reports the comparison of the energy per bit, i.e., the energy spent to transport a bit of information from the source to the destination. Both energy-efficient networks and a standard network are considered. The energy per bit is computed as[11]

$$\mathrm{EB}(t) = \frac{\mathcal{P}_{\mathrm{tot}}(t)}{\sum\limits_{sd} t^{sd}(t)}. \tag{24}$$

The plot shows that a green network design allows to reduce the cost of transporting information during the whole day, with higher gains during low traffic periods, i.e., when the amount of spare resources allows to turn off a large number of devices. Also in this case, the best savings are obtained using the OPT-V3 algorithm and the MP-MP heuristic, whose lines in the figure are almost overlapping. Interestingly, during off-peak hours, the energy per bit is higher than during peak hours.

To compute the total saving per year, Table III shows the yearly energy saving obtained using the optimal formulation OPT-V3 and the MP-MP heuristic with the consumption figures of Table I. Note that the daily energy consumption $\mathrm{ED}_{\mathrm{tot}}$ has been computed from $\mathcal{P}_{\mathrm{tot}}(t)$ by

$$\mathrm{ED}_{\mathrm{tot}} = \int_{t=0}^{T} \mathcal{P}_{\mathrm{tot}}(t)dt \approx \sum_{i=0}^{N-1} \mathcal{P}_{\mathrm{tot}}[t_i]\Delta t_i. \tag{25}$$

---

[10] This overprovisioning is a typical ISP policy. This is due to the low cost of high-bandwidth technologies in the metro area and allows to accommodate unexpected future increases in the number of customers.

[11] Notice that $1\ \mathrm{W} = 1$ J/s, so that the measurement unit of energy per bit is J/b.

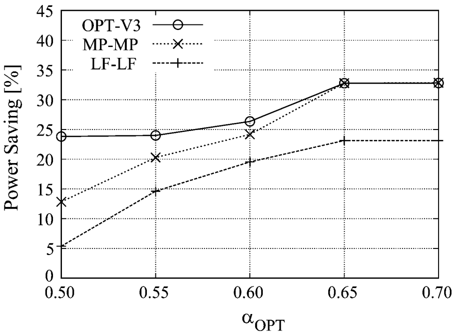

Fig. 17. Power saving versus α.

Notice that since $T = 24$ h, 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 $\mathrm{EY}_{\mathrm{tot}} = \mathrm{ED}_{\mathrm{tot}} \times 365$, and then we can compute the energy savings as

$$\text{Energy Saving} = 1 - \frac{\mathrm{EY}_{\mathrm{tot}}^{\mathrm{GREEN}}}{\mathrm{EY}_{\mathrm{tot}}^{\mathrm{STD}}} \tag{26}$$

where $\mathrm{EY}_{\mathrm{tot}}^{\mathrm{STD}}$ and $\mathrm{EY}_{\mathrm{tot}}^{\mathrm{GREEN}}$ are the yearly energy consumed by a standard network and a green network, respectively.

We can see that with a green approach it is possible to save more than 3.7 GWh in a year, corresponding to a percentage saving of more than 29% and a monetary saving of more than 340 000 Euros per year considering current electricity prices of 0.092 Euro/kWh. Moreover, using the MP-MP heuristic, the savings are only 2.4% away from the optimal solution.

2) Impact of the Overprovisioning Factor: In the last set of experiments, we performed a study on the impact of the QoS parameter α to observe its impact on the effectiveness of the green approach. We assumed a network design that considers the QoS constraint of $\alpha = 0.5$ during peak-hour demand. What happens if now we relax the constraint and raise the maximum load factor to values larger than 0.5?[12]

The rationale of changing α is to analyze the impact of our approach as a function of the maximum utilization that is enforced to guarantee the QoS. We therefore consider peak-hour traffic and evaluate the energy saving that can be achieved by the network for different values of α. For the sake of simplicity, we report only the average results in the rest of this section.

Fig. 17 shows the results of increasing the maximum load on the links to $\alpha \in [0.5, 0.7]$. All algorithms show improvements for α up to 0.65. The intuition is that by allowing more traffic on links, it is possible to actually turn off more nodes and links achieving higher savings. For $\alpha > 0.7$ no improvement is observed since the connectivity constraint limits the number of nodes and links that can be powered off. In particular, the OPT-V3 algorithm is able to save an additional 8% of power by increasing the tolerated offered load from $\alpha = 0.5$ to $\alpha = 0.65$. Considering the heuristics, MP-MP is close to the optimal solution for $\alpha > 0.6$, while LF-LF performs consistently worse.

VI. IMPLEMENTATION DISCUSSION

While the results in this work show that there is a great opportunity to save energy consumption in a real network, today's technology does not fully support the selective shutdown of links and nodes. First, support to equipment 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 controlled and simultaneous "failures" of nodes and links. Notice, however, that contrary to failures that are unpredictable, the decision to power off a device can be advertised in advance, so a smooth transient can be adopted. Note that traffic engineering approaches can be integrated. Furthermore, new primitives to existing routing protocols may be introduced, including the notification of topology changes when multiple devices are powered off and the signaling of the power state. These issues are, however, outside the scope of this paper.

In our work, we assume that the energy-aware algorithms are performed by a centralized controller that selects the subset of resources that must be powered on to meet the current traffic demand. In this context, a tradeoff among responsiveness to traffic variation and time to perform the algorithms emerges. We foresee a deployment scenario in which the network configuration changes slowly, e.g., every 30 min. This is compatible with the slow and daily variation of traffic in current backbone networks. Guard thresholds can be easily integrated to take into account possible sudden increase of traffic (e.g., triggering a quick power-on of devices). This alleviates the constraint due to the algorithm complexity since a new solution must be computed within minutes. Note that having few reconfigurations per day mitigates the eventual extra energy cost and additional latencies that devices may introduce when changing power state (e.g., energy spent to save/recover router state on/from disk, latencies in settling the long-haul amplifiers, etc.).

Finally, a proper discussion on what an operator needs to sacrifice to achieve good energy savings is in place. Clearly, changing network configuration impacts recovery mechanisms and flexibility in traffic engineering. Notice that it could be possible to include more constraints, e.g., explicitly taking into account both the presence of k-alternative paths and of more complex routing schemes. This will increase the complexity of the formulation, questioning the finding of optimal solutions. Still, heuristic complexity would be similar since the only change would be a slightly more complex constraint to check. Second, considering reliability constraints, we envisage the adoption of technologies that allow to quickly power on devices. Their adoption is in line with our proposal.

VII. RELATED WORK

The study of power-saving network devices has been introduced in recent years, starting from the pioneering work of [9]. In [17], the authors describe a power benchmarking scheme for network devices and propose an energy index to compare different network devices such as routers, switches, and access points (APs). In [18], the ideas of adaptive link rate (ALR) and protocol proxying are proposed. Both these techniques require to change protocols, and both consider pairs of devices, e.g.,

---

[12] Notice that during the computation of the traffic matrix, we set $\alpha = 0.5$ also in this case.

routers, sharing the same link or a pair of high/low performance servers.

More recently, some effort was devoted to investigate how to reduce the power consumption of the entire network infrastructure, and not of single or few components only. In [10], measurements of power consumption of networking devices are first presented, then the authors consider a network topology and evaluate the total network consumption given the power footprint of each element. They consider two scenarios: In the first one, all devices are turned on, while in the second one, only the minimum number of elements to guarantee the service are actually powered on. The reduction of the corresponding power consumption is finally evaluated on simple network topologies. Differently from our work, the authors use a standard formulation of the CMCF problem, showing that the complexity grows very fast as the number of devices increases, making the problem very expensive to be solved even for small networks. In [11] and [12], we have proposed a preliminary study of the possible savings obtained with the heuristics considered in this paper, in which we further compared the heuristics to the optimal formulations, considering also the impact of computation times. In [13], the authors propose a simple heuristic to shut down links when bundles of multiple physical cables are present. This solution can be potentially integrated with our approach to further increase the energy savings for networks with bundled links. In [19], the authors solve an energy-aware routing problem considering different technology assumptions. Differently from our work, the authors consider less complex topologies composed by core nodes only, which are all sources and sinks of traffic so that only links can be powered off.

The potential of energy-aware routing algorithms in backbone networks is presented in [20]. In particular, the authors propose a modification of the OSPF protocol to switch off links in an IP network, showing that more than 60% of links can be potentially turned off in a realistic network. This work encourages us to investigate a possible distributed solution for our approach.

The evaluation of power-aware management schemes for networks is argued in [21]. In particular, the authors evaluate the energy savings from sleeping, i.e., during absence of packets, and the possible savings from rate adaption. This is a complementary approach with respect to our work. These power adaptation schemes can help to further reduce the power under light load conditions, such as the very lightly loaded links during nighttime (see Fig. 13).

In [22], the authors exploit the idea of exchanging energy profiles among devices to reduce the overall power consumption during routing and traffic-engineering operations. Differently from our work, the authors do not take into account the traffic variation, and the topology considered is composed by only one level of nodes. Moreover, in our work the impact of QoS constraints is investigated.

The capacitated network design (CMCF) problem is well known in the literature (see [23] for an overview). In the past years, the primary goal of the CMCF problem was to maximize the utilization of network devices or the quality of service perceived by the users. In this paper, we propose the novel idea of maximizing the power savings, imposing the QoS in the formulation as a constraint.

Considering the single devices, energy-aware solutions involving switches and software routers have been proposed in [24] and [25], respectively.

In [26], authors have proposed efficient solutions to reduce the energy consumption of switches and links in data-center networks. Differently from our solution, the authors consider more regular topologies. However, the obtained power savings are comparable with our results.

Finally, energy-aware data centers taking into account even the variations of the hourly electricity prices have been studied (see, e.g., [27]). This optimization could be fruitfully adopted by a telecommunication operator to further improve the power savings.

VIII. Conclusion and Future Work

In this paper, we faced the problem of reducing the power consumption of backbone networks. Our aim was to find the minimal set of routers and links to satisfy a given traffic demand under connectivity and quality-of-service constraints. We have first formulated the problem using an ILP formulation, showing that the problem falls in the class of multicommodity flow problems and therefore it is NP-complete. Computationally efficient formulations have been introduced, adding additional constraints and reduced notations. Simple heuristics have been proposed to reduce the computation time in case large networks are considered.

Both optimal formulations and heuristics have been extensively compared on synthetic and real topologies, taking into account real traffic figures and considering the daily variation of traffic. Results show that the possible energy savings can be large, especially during off-peak times, when traffic is low.

These encouraging results are supporting the effort spent by the research community. In particular, they call for the availability of both devices that support different power states and the design of distributed algorithms to allow the real implementation of an energy-aware control plane.

References

[1] A. P. Sokolov, P. H. Stone, C. E. Forest, R. Prinn, M. C. Sarofim, M. Webster, S. Paltsev, C. A. Schlosser, D. Kicklighter, S. Dutkiewicz, J. Reilly, C. Wang, B. Felzer, J. Melillo, and H. D. Jacoby, "Probabilistic forecast for 21st century climate based on uncertainties in emissions (without Policy) and climate parameters," MIT, Cambridge, MA, MIT Global Change Program Report, Jan. 2009.

[2] "An inefficient truth," Global Action Plan, London, U.K., Dec. 2007 [Online]. Available: http://www.globalactionplan.org.uk/

[3] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt, and P. Demeester, "Worldwide energy needs for ICT: The rise of power-aware networking," in *Proc. IEEE ANTS*, Bombay, India, Dec. 2008, pp. 1–3.

[4] "SMART 2020 report, Enabling the low Carbon Economy in the information age," The Climate Group, 2008 [Online]. Available: http://www.theclimategroup.org

[5] S. Pileri, "Energy and communication: Engine of the human progress," presented at the INTELEC Rome, Italy, 2007 [Online]. Available: http://www.ega.it/intelec2007/img/Stefano_Pileri.pdf

[6] "Regulatory authority for electricity and gas," Italian Regulatory Authority for Electricity and Gas, Milan, Italy, 2010 [Online]. Available: http://www.autorita.energia.it/

[7] M. Etoh, T. Ohya, and Y. Nakayama, "Energy consumption issues on mobile network systems," in *Proc. IEEE SAINT*, Turku, Finland, Jul. 2008, pp. 365–368.

[8] R. Tucker, J. Baliga, R. Ayre, K. Hinton, and W. Sorin, "Energy consumption in IP networks," in *Proc. ECOC Symp. Green ICT*, Brussels, Belgium, 2008, p. 1.

[9] M. Gupta and S. Singh, "Greening of the Internet," in *Proc. ACM SIGCOMM*, Karlsruhe, Germany, Aug. 2003, pp. 19–26.

[10] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," in *Proc. IEEE INFOCOM*, Phoenix, AZ, Apr. 2008, pp. 457–465.

[11] L. Chiaraviglio, M. Mellia, and F. Neri, "Reducing power consumption in backbone networks," in *Proc. IEEE ICC* , Dresden, Germany, Jun. 2009, pp. 1–6.

[12] L. Chiaraviglio, M. Mellia, and F. Neri, "Energy-aware backbone networks: A case-study," in *Proc. 1st GreenComm*, Dresden, Germany, Jun. 2009, pp. 1–5.

[13] W. Fisher, M. Suchara, and J. Rexford, "Greening backbone networks: Reducing energy consumption by shutting off cables in bundled links," in *Proc. 1st ACM SIGCOMM Workshop Green Netw.*, New Delhi, India, Aug. 2010, pp. 29–34.

[14] T. G. Crainic, M. Gendreau, and I. Ghamlouche, "Cycle-based neighbourhoods for fixed-charge capacitated multicommodity network design, Centre de Recherche sur les Transports, Universite de Montreal, Montreal, QC, Canada, Tech. Rep. CRT-2001–01, 2001.

[15] L. Fratta, M. Gerla, and L. Kleinrock, "The flow deviation method: An approach to store-and-forward network design," *Networks*, vol. 3, no. 2, pp. 97–133, 1973.

[16] F. K. Hwang, D. S. Richards, and P. Winter, "The Steiner tree problem," *Ann. Discr. Math.*, vol. 53, pp. 1–335, 1992.

[17] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "A power benchmarking framework for networking devices," in *Proc. IFIP Netw. Conf.*, Aachen, Germany, May 2009, pp. 795–808.

[18] K. Christensen, C. Gunaratne, and B. Nordman, "Managing energy consumption costs in desktop PCs and LAN switches with proxying, split TCP connections, and scaling of link speed," *Int. J. Netw. Manage.*, vol. 15, no. 5, pp. 297–310, Sep. 2005.

[19] A. P. Bianzino, C. Chaudet, F. Larroca, D. Rossi, and J. L. Rougier, "Energy-aware routing: A reality check," in *Proc. 3rd GreenComm*, Miami, FL, Dec. 2010, pp. 1422–1427.

[20] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, "An energy saving routing algorithm for a green OSPF protocol," in *Proc. IEEE INFOCOM Workshops*, San Diego, CA, Mar. 2010, pp. 1–5.

[21] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing network energy consumption via sleeping and rate-adaptation," in *Proc. USENIX/ACM NSDI*, San Francisco, CA, Apr. 2008, pp. 323–336.

[22] J. Restrepo, C. Gruber, and C. Machoca, "Energy profile aware routing," in *Proc. 1st GreenComm*, Dresden, Germany, Jun. 2009, pp. 1–5.

[23] M. Pioro and D. Medhi*, Routing, Flow and Capacity Design in Communication and Computer Networks*, ser. Networking. New York: Morgan Kaufmann, 2004.

[24] G. Ananthanarayanan and R. H. Katz, "Greening the switch," in *Proc. HotPower*, San Diego, CA, Dec. 2008.

[25] R. Bolla, R. Bruschi, and A. Ranieri, "Green support for PC-based software router: Performance evaluation and modeling," in *Proc. IEEE ICC*, Dresden, Germany, Jun. 2009, pp. 1–6.

[26] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, "ElasticTree: Saving energy in data center networks," in *Proc. NSDI* , San Jose, CA, Apr. 2010, p. 17.

[27] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the electric bill for Internet-scale systems," in *Proc. ACM SIGCOMM*, Barcelona, Spain, Aug. 2009, pp. 123–134.

Luca Chiaraviglio (M'09) received the Ph.D. degree in electronics and communication engineering from the Politecnico di Torino, Turin, Italy, in February 2011.

Between August 2009 and April 2010, he was a visiting scholar with the Web and InterNetworking Group (WING), Boston University, Boston, MA, under the supervision of Prof. Ibrahim Matta. His main research interests are in the field of energy-aware design for backbone and access networks.

Marco Mellia (M'97–SM'08) received the Ph.D. degree in telecommunications engineering from the Politecnico di Torino, Turin, Italy, in 2001.

In 1999, he was with the Computer Science Department, Carnegie Mellon University, Pittsburgh, PA. Since April 2001, he has been with the Electrical Engineering Department, Politecnico di Torino. He has coauthored over 150 papers published in international journals and conferences. His research interests are in the fields of traffic measurement, P2P applications, and energy-aware network design.

Dr. Mellia participated in the program committees of several conferences including IEEE INFOCOM and ACM SIGCOMM.

Fabio Neri (M'98–SM'08) was born in Novara, Italy, in 1958. He received the M.S. and Ph.D. degrees in electrical engineering from the Politecnico di Torino, Turin, Italy, in 1981 and 1987, respectively.

He was a Full Professor with the Electronics Department, Politecnico di Torino, where he led research activities on optical networks and on switching architectures. He coordinated the participation of his research group into several Italian research projects. He participated in a number of European projects on WDM networks, and was the coordinator of the FP6 Network of Excellence (NoE) e-Photon/ONe on optical networks, which involved 40 European institutions. He was project leader of the FP7 NoE TREND on energy efficiency in telecommunication networks. He coauthored over 200 papers published in international journals and presented in leading international conferences. His research interests were in the fields of performance evaluation of communication networks, high-speed and all-optical networks, packet switching architectures, and queuing theory.

Prof. Neri served on the Editorial Board of the IEEE/ACM TRANSACTIONS ON NETWORKING and was Co-Editor-in-Chief of the *Optical Switching and Networking Journal*. He participated in the technical program committees of several conferences, including IEEE INFOCOM and IEEE GLOBECOM. He was the General Co-Chairman of ECOC 2010.