Energy Management in IP Traffic Engineering with Shortest Path Routing

Edoardo Amaldi, Antonio Capone, Luca G. Gianoli, Luca Mascetti
DEI - Politecnico di Milano, Italy

Abstract—**Even if the Internet is commonly considered a formidable mean to reduce the impact of human activities on the environment, its energy consumption is rapidly becoming an issue due to the exponential traffic growth and the rapid expansion of communication infrastructures worldwide. One of the first options available to reduce energy consumption is that of limiting wastes. Actually, even if it has been shown that traffic load greatly varies over time and rarely saturates available capacity, the energy consumed by the network is almost constant as if it were fully loaded at all time. In this paper we propose an IP traffic engineering approach that allows to adapt the network energy consumption to different daily traffic scenarios switching on and off communication interfaces (links) and entire routers. We focus on routing domains of the Internet Service Providers where the popular OSPF (Open Shortest Path First) protocol is adopted and we consider the problem of switching off network elements (links and routers) and of adjusting link weights so as to minimize the energy consumption and the network congestion. We present two heuristics for this problem: the Greedy Algorithm for Energy Saving (GA-ES) and the Two-stage Algorithm for Energy Saving (TA-ES). The computational results for three real network topologies show that it is possible to switch off up to 80% of the core nodes during low traffic periods (night hours), and up to 65% during moderate traffic periods, while guaranteeing the same point-to-point service quality, and moderately increasing the network congestion.**

I. INTRODUCTION

In the last years, the exponential growth of Internet has been accompanied by a considerable increase in its energy consumption. Data reported in [1] show that in 2007, Internet had been responsible for 5.5% of the total energy consumption in the world and that the annual increment rate can be estimated around 20-25%. For these reasons, the issues of energy saving in IP networks and of power awareness in network design have recently become of great interest in the scientific community and have attracted the interest of device manufacturers and Internet Service Providers (ISP) [2].

Since the seminal work by Gupta and Singh [3], the research community has started developing technologies for manufacturing energy efficient network devices, methodologies for power aware network design, as well as energy management strategies for reducing energy wastes of networks in operation [4], [5]. In this paper we focus on energy management of the network and its elements.

Energy management in the Internet exploits the fact that networks are designed and dimensioned to serve the estimated peak traffic demand. Usually, during network operation, traffic load varies remarkably over time and even during peak hours it is usually well below network capacity. Unfortunately, current network device architectures and transmission technologies make their power consumption almost independent of the traffic load.

A possible approach to improve the Internet energy performance is to optimize the operation of individual routers in order to adjust their consumption according to the traffic load [6]. However, a coordinated management of the energy consumption of the network that is able to dynamically switch off and on links and nodes according to traffic variations can achieve a better performance than the management of individual devices only.

In order to reduce energy consumption by switching off part of the network, we need to guarantee that the remaining portion of the network has enough capacity to support the traffic demands and we have to reroute traffic through active routers. The routing protocol is thus a key element in the energy management and it may limit possible strategies.

We consider the most widely used Internal Gateway Protocol (IGP) in IP networks, namely the Open Shortest Path First (OSPF) protocol. The link state approach adopted by OSPF to distribute topology information and the local calculation of shortest paths based on a set of link weights prevent per-flow routing optimization as for instance in Multi-Protocol Label Switching (MPLS). Therefore, traffic engineering in OSPF is forced to optimize the set of shortest paths adjusting the link weights [7].

In this paper we propose a traffic engineering strategy that aims at minimizing the energy consumption and the network congestion, while guaranteeing that all the traffic demands are routed. We present and compare two heuristic algorithms, GA-ES (Greedy Algorithm for Energy Saving) and TA-ES (Two-stage algorithm for Energy Saving), for tackling the power-aware OSPF link weights configuration problem. Both GA-ES and TA-ES exploit in different ways the IGP-WO algorithm originally proposed in [7] that, given a network topology and a traffic matrix, aims at a set of OSPF link weights that minimizes the network congestion.

The paper is organized as follows. In Section II we summarize related work. In Section III we describe the problem and present a mathematical programming formulation. Since solving it to optimality even for small size instances requires a very high computing time, in Sections IV and V we propose, respectively, GA-ES and TA-ES. In Section VI we report an discuss the computational results obtained on three real network topologies with constant and poisson traffic matrices. The results indicate that it is possible to switching off up to

Corresponding author: Antonio Capone, email: antonio.capone@polimi.it.

80% of the core nodes during low workload periods (night hours), and up to 65% during moderate workload periods, while guaranteeing the same point-to-point service quality and moderately increasing the network congestion. Finally, Section VII contains some concluding remarks.

II. RELATED WORK

As discussed in [5], the problem of energy consumption in communication networks can be tackled at different levels. Part of the research aims at developing a new generation of network devices with a better energy performance [6], also through adaptive link rate (ALR) techniques that adjust link capacities according to traffic needs [8], [9].

Other works study models of router power consumption, in order to define network design strategies that can guarantee at the same time low energy consumption, good performance and resilience to failures [10].

Our work, instead, is focused on efficient models and algorithms for a power-aware management of IP networks through routing strategies that select a subset of network resources able to support the traffic load. As far as this green IP routing research area is concerned, only a few works have appeared recently.

In [11] the authors describe some heuristics that, given a traffic matrix and a fully powered network, are able to switch off nodes and links while respecting traffic constraints.
In [12] the researchers present some Energy-Aware Traffic Engineering (EATe) techniques, that aim at optimizing links and routers power consumption, by considering their rate-dependant energy profiles. Assuming the network elements consuming different amount of energy when working at different fixed rates, EATe algorithms try to switch off the underutilized elements, or in alternative, to reduce their rate, by reallocating part of the traffic in other network portions without forcing any element to increase its current rate. Unlike in our work, the approaches in [11] and [12] follow a flow-based strategy that is suitable for the Internet geographic backbone where label switching protocols are adopted and not for routing domains based on shortest paths.

The energy management algorithm for IP networks called Energy Aware Routing (EAR) Algorithm and presented in [13] is able to switch off network elements exploiting a modified version of the OSPF protocol. EAR algorithm selects a subset of routers, Importers Router (IR), that do not calculate their own shortest path tree (SPT) but use that of some neighboring routers, Exporters Routers (ER). In general a small number of active SPTs reduces also the number of links used that can be switched off. There are several important differences between our work and [13]. As in the literature on traffic engineering in OSPF [14], we keep the OSPF protocol unchanged and focus on optimizing the link weights, while in the EAR algorithm the weights are assumed to be given and the protocol needs to be modified to implement ERs and IRs. Moreover, we explicitly consider link capacity limitations and minimize the network congestion level in order to guarantee service quality, while in [13] neither traffic load nor network capacity are considered.

III. PROBLEM AND MATHEMATICAL PROGRAMMING FORMULATION

OSPF is the most commonly adopted intra-domain internet routing protocol. Traffic demands are routed along the shortest paths, and are evenly split at nodes where more outgoing links belong to shortest paths to the destination. Link weights are managed by network operators, who may modify them in order to optimize routing and reduce network congestion.

Let the directed graph $G = (V, A)$ represent the network topology, where V is the set of nodes and A the set of links. We distinguish two types of nodes: *edge* nodes and *core* nodes. Edge nodes can be both source and destination of traffic demands, while core nodes play only the role of transit routers. Let D denote the traffic matrix, where d_{ij} is the traffic demand for each pair of edge nodes i and j, and $d_{ij} = 0$ for all other pairs of nodes.

We consider the following extension of the IGP weight optimization problem for intra-domain Traffic Engineering [15], that we refer to as *Green Traffic Engineering* (GTE).

> **GTE**: Given a directed graph $G = (V, A)$, representing the topology of an IP network composed by routers and links, and a traffic matrix D, decide which network elements (routers and links) to switch off and determine the OSPF link weights so as to minimize the total network energy consumption (primary objective) and the network congestion (secondary objective), while guaranteeing that all the traffic demands are routed and all link capacities are satisfied.

According to the distinction between primary and secondary objectives, we first look for a sub-network with minimum total energy and then minimize the sub-network congestion. The secondary objective can be minimized by adjusting the OSPF link weights of the sub-network corresponding to the active elements. The unnecessary routers and links can then be excluded from the shortest path trees, by assigning a very large value to the corresponding link weights.

Let p_{ij} and p_k be the power consumption of link (i, j) and node k respectively. Let c_{ij} be the capacity of the link (i, j). If the binary decision variables x_{ij} and y_k represent the power status (on/off) of links and routers respectively, a mathematical programming formulation of the GTE problem is given by:

$$\min \sum_{(i.j)\in A} p_{ij}\, x_{ij} + \sum_{k\in V} p_k\, y_k \tag{1}$$

s. t.

$$x_{ij} \le y_i \qquad (i,j) \in A \tag{2}$$

$$x_{ij} \le y_j \quad (i,j) \in A \tag{3}$$

$$\sum_{i\in V} f^t_{it} = \sum_{s\in V} d_{st} \quad t \in N \tag{4}$$

$$\sum_{j\in V} f^t_{vj} - \sum_{i\in V} f^t_{iv} = d_{vt} \quad v, t \in V, t \neq v \tag{5}$$

$$\sum_{t\in V} f^t_{ij} \le x_{ij}\alpha c_{ij} \quad (i,j) \in A \tag{6}$$

$$0 \le z_i^t - f_{ij}^t \le (1 - u_{ij}^t) \sum_{v \in V} d_{vt} \quad t \in V, (i,j) \in A \tag{7}$$

$$f_{ij}^t \le u_{ij}^t \sum_{v \in V} d_{vt} \quad t \in V, (i,j) \in A \tag{8}$$

$$0 \le r_j^t + \omega_{ij} - r_i^t \le (1 - u_{ij}^t) M \quad t \in V, (i,j) \in A \tag{9}$$

$$1 - u_{ij}^t \le r_j^t + \omega_{ij} - r_i^t \quad t \in V, (i,j) \in A \tag{10}$$

$$u_{ij}^t \le x_{ij} \quad (i,j) \in A, t \in V \tag{11}$$

$$1 \le \omega_{ij} \le \omega_{max} \quad (i,j) \in A \tag{12}$$

$$\omega_{ij} \in Z \quad (i,j) \in A \tag{13}$$

$$u_{ij}^t \in \{0,1\} \quad t \in V, (i,j) \in A \tag{14}$$

$$x_{ij}, y_k \in \{0,1\} \quad (i,j) \in A, k \in V \tag{15}$$

$$f_{ij}^t \ge 0 \quad (i,j) \in A, t \in V \tag{16}$$

$$r_i^t, z_i^t \ge 0 \quad i, t \in V, \tag{17}$$

where M is a large enough constant. The objective function (1) aims at minimizing the total power consumption of the network. Constraints (2-3) ensure that if a node is switched off all incident links are turned off. Obviously a node can be switched off only if there are no traffic demands having it as source or destination (edge or core node). Constraints (4-5) are the classical flow conservation constraints, where the (real) positive variable f_{ij}^t indicates the amount of flow routed through the link $(ij) \in A$ destined to node $t \in N$. Constraints (6) are the max-utilization constraints imposing that the total flow through each link does not exceed the link maximum utilization and forcing the flow to 0 if the link (i,j) is powered off; the parameter α is comprised between 0 and 1. The binary variables $u_{ij}^t = 1$ appearing in Constraints (7-9) describe the routing configuration: $u_{ij}^t = 1$ if and only if the link (i,j) belongs to one of the shortest paths from node i to node t. Constraints (7) make sure that if $u_{ij}^t = 1$ then the flow f_{ij}^t destined to node t is equal to the (real) variable z_i^t, which is the common value of the flow assiggned to all links outgoing from i and belonging to the shortest paths from i to t. Constraints (8) force $f_{ij}^t = 0$ for all links (i,j) that do not belong to a shortest path to node t. Finally, the shortest path routing Constraints (9-14) assure that the routing vector u defines shortest paths consistent with the link weight vector ω and forbid switched off links to belong to a shortest path. For each pair of nodes j and t, the (real) variable r_j^t corresponds to the length of the shortest path from node j to node t.

Unfortunately, the above mixed integer programming formulation turns out to be very challenging even for small size networks. Indeed, the subproblem of adjusting the link weights so as to minimize the network congestion is already very heavy computationally. For instance, in [16] computing times larger than 69 minutes are reported to solve with CPLEX the subproblem instances with only 6 nodes and 28 links.

	LF	**LL**	**SW**
LF	LF-LF	LL-LF	SW-LF
TE	LF-TE	LL-TE	SW-TE

Table I: Combinations of sorting policies, the rows correspond to link criterion and the columns to router criterion.

IV. GREEDY ALGORITHM FOR ENERGY SAVING (GA-ES)

A first simple approach to tackle the Green Traffic Engineering problem is to adopt a greedy strategy. Given a network topology G and a traffic matrix D, the idea is to sort the network elements (links and routers) according to some criteria, to consider them in that order and to try to switch off as many of them as possible.

In the initial phase of our Greedy Algorithm for Energy Saving (GA-ES) a set of values of the OSPF link weights is computed with the IGP-WO algorithm proposed in [7] by Fortz and Thorup. The IGP-WO algorithm is a local search heuristic looking for a set of link weights that minimizes an objective function, the overall cost of link utilization, which is directly related to the network congestion. IGP-WO tries to balance the load in the given network topology.

GA-ES is based on a sorting criterion $g(e)$, where e is a network element, consisting of two policies, one for the routers (nodes) and one for the links. We use very intuitive criteria for sorting network elements, three for sorting routers and two for sorting links. Two of these policies, Least-Link (LL) and Least-Flow (LF), were already used in a different context in [11]. In LL and LF the routers are sorted in a non-decreasing order according to, respectively, the router degree (number of incident links) and the total amount of traffic flowing through the router.

We also propose two policies for sorting network elements based on the value of the OSPF link weights. In Traffic-Engineering (TE) all the links are sorted by non-increasing values of the link weights. In Sum-of-Weights (SW) the routers are sorted by non-increasing values of the sum of all their active link weights. The six combinations of policies for sorting routers and links that we consider are summarized in Table I.

Once the initial set of OSPF link weights is determined with IGP-WO for the initial network topology, all the network elements are sorted according to the two selected criteria. First all nodes are considered (because of their higher energy consumption) and then all links. At each step of GA-ES we verify whether the first available element can be turned off. The considered element is actually turned off if the OSPF routing determined on the reduced network by the weights of the active links, is able to support the traffic matrix (all traffic demands) without exceeding the links maximum utilization $U_{MAX} = \alpha$ ($0 < \alpha \le 100\%$). With $\alpha = 100\%$ the algorithm is able to maximally exploit network resources, while with $\alpha < 100\%$, a part of the network capacity remains unused, being thus able to absorb the unpredictable variations of traffic. After turning off any element, the order of the remaining network elements is updated. GA-ES proceeds until all the network elements have been tested.

It is worth pointing out that we have adopted a natural

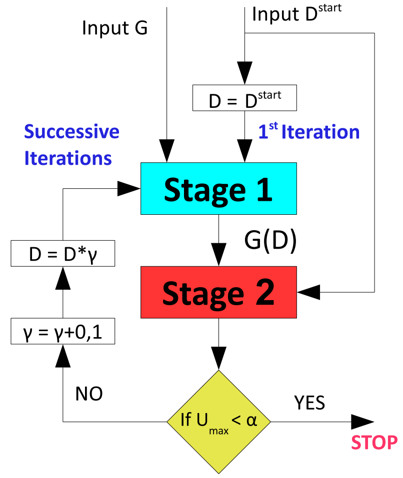

Figure 1: TA-ES flow chart.

shutting down strategy rather than a constructive strategy. Starting from a traffic matrix that is routable in the initial network, GA-ES always works on sub-networks that are able to support all the traffic demands. Note that the set of values of the link weights is not determined at each iteration, but only twice: at the beginning for the initial network topology and at the end of the algorithm for the resulting sub-network.

V. TWO-STAGE ALGORITHM FOR ENERGY SAVING (TA-ES)

Since the problem of finding a minimum energy sub-network of G with an optimized set of OSPF link weights is very challenging, we split it into two distinct stages: the switching-off stage and the feasible routing stage. Figure 1 shows the complete flow chart of the procedure.

The first stage, which receives as input the network topology G and the given traffic matrix denoted by D, aims at selecting the set of network elements that could be switched off. This is achieved by solving the following Integer Linear Program (ILP), that is a subset of the GTE formulation (1-16), with a gap $\leq 1\%$:

$$\min \quad \sum_{(i.j)\in A} x_{ij} + \sum_{k\in V} \lceil \frac{3g_k}{2} \rceil y_k \tag{18}$$

s. t.

$$x_{ij} \leq y_i \qquad (i,j) \in A \tag{19}$$

$$x_{ij} \leq y_j \quad (i,j) \in A \tag{20}$$

$$\sum_{i\in V} f^t_{it} = \sum_{s\in V} d_{st} \quad t \in N \tag{21}$$

$$\sum_{j\in V} f^t_{vj} - \sum_{i\in V} f^t_{iv} = d_{vt} \quad v,t \in V, t \neq v \tag{22}$$

$$\sum_{t\in V} f^t_{ij} \leq x_{ij}\alpha c_{ij} \quad (i,j) \in A \tag{23}$$

$$x_{ij}, y_k \in \{0,1\} \quad (i,j) \in A,\, k \in V \tag{24}$$

$$f^t_{ij} \geq 0 \quad (i,j) \;\in A,\, t \in V. \tag{25}$$

Network	Nodes	Links	Edge$_{node}$	Core$_{node}$	%Core$_{node}$
Telstra	104	302	65	39	37.5
Ebone	87	322	31	56	64.4
Exodus	79	294	38	41	51.9

Table II: Rocketfuel network topologies.

The objective function (18) aims at minimizing the network power consumption. Compared to (1), in (18) we take $p_{ij} = 1$ for every link (i,j) (we assume that the links consumption is unitary) and $p_k = \lceil \frac{3g_k}{2} \rceil$ for every node k, where g_k is the degree of node $k \in V$. Considering the set composed by a router and its links, the router is then responsible for about 66% of the total power consumption. Thus the formulation favours switching-off nodes with higher degree. Constraints (19-25) are identical to Constraints (2-6) and (15-16) of the GTE formulation. Note that traffic demands routing is considered fully splittable, see Constraints (25).

The formulation used in the first stage falls within the well-known class of capacitated multi-commodity minimum cost flow problems (CMCF) [17]. Although CMCF problems are NP-hard, a solver like CPLEX can provide feasible solutions with an acceptable gap ($\simeq 1\%$) with respect to the best bound found for instances of about 100 nodes and 300 links in a few hours of computing time.

The feasible routing stage, which receives as input the sub-network $G(D)$ determined at the switching-off stage, aims at finding a set of OSPF link weights that allows to route through the sub-network all the traffic demands according to shortest paths, without exceeding the links maximum utilization $U_{MAX} = \alpha$ $(0 < \alpha \leq 100\%)$. The OSPF link weights configuration of the second stage is determined with the IGP-WO algorithm mentioned in Section IV. Unfortunately, given the sub-network topology $G(D)$ obtained at the first stage and its respective traffic matrix D, there is no guarantee that there exists a set of OSPF link weights allowing feasible routing of all traffic demands. This may occur because the first stage considers a fully splittable routing that can be hardly reproduced by the OSPF protocol. The IGP-WO algorithm is repeated with five different random sets of initial OSPF links weights; in case no feasible OPSF weights set is found ($U_{MAX} > \alpha$), we proceed as shown in Figure (1). We slightly increase the original traffic matrix D by multiplying it for a fixed parameter γ, and repeat the first stage using the new increased traffic matrix D. γ is equal to 1 at the first iteration, and is increased by $0,1$ every time the second stage fails to find a feasible set of weights. This operation clearly leads to a sub-network $G(D)$ with more active elements as input of the second stage (the second stage is always run with the original traffic matrix D^{start}). The TA-ES algorithm terminates when a feasible set of link weights is found at the second stage.

VI. COMPUTATIONAL RESULTS

A. Network Topologies and Traffic Matrices

We have carried out computational tests on the real network topologies provided by the Rocketfuel project [18]. In particular, we have worked with the three networks presented in Table

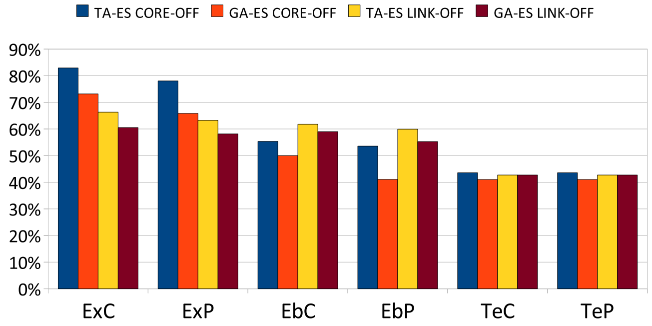

Figure 2: % of core routers and links switched off by GA-ES and TA-ES for the constant and poisson traffic matrices.

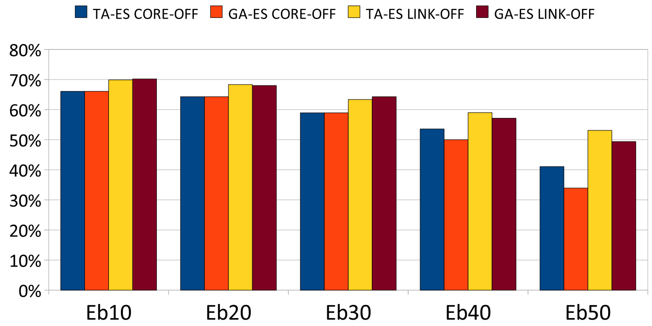

Figure 3: % of core routers and links switched off by GA-ES and TA-ES for the LP Ebone scaled matrices.

II. Unfortunately, although the network topologies are known, no accurate information is available concerning link capacities. Therefore we have assumed that all the network links have the same capacity. Also information about the subdivision between edge and core nodes is missing. Since edges nodes can be both source and destination of traffic demands and core nodes play only a role of transit routers, core nodes are the only one that can be powered off. We have randomly selected a set of edge routers for each of the three topologies. To avoid the trivial and unrealistic case where core leaf nodes can be easily powered off, all the leaf nodes are considered as edge nodes. At least one edge node has also been selected for each city.

Moreover, in our tests we have for simplicity considered a max utilization limit of 100%.

The traffic matrices have been generated in two ways:

1) **Constant and Poisson**: generated using the Totem toolbox [19], the maximum load matrices with constant and poisson traffic distribution that can be supported by the (complete) networks with OSPF *hop-count* routing.
2) **LP formulation matrices:** generated by scaling with a parameter $\alpha \in (0, 1)$ a maximum supported traffic matrix obtained with a linear programming (LP) formulation. By maximum supported traffic matrix, we mean that all the traffic demands can only be satisfied by switching on all the network links/routers and performing fully splittable routing.

B. Tests

The computational results reported in Figures 2, 3, and Table III, show that it is possible to switch off up to 80% of the core nodes during low workload periods (night hours), and up to 65% during moderate workload periods. In case of extremely low workload conditions (Ex10-Ex20-Eb10-Eb20-Te10-Te20), GA-ES and TA-ES achieve the same results in terms of percentage of network elements switched off, while GA-ES shows better performance than TA-ES in terms of congestion level. Instead, in case of moderate workload conditions (Ex30-Ex40-Ex50-Eb30-Eb40-Eb50-ExC-ExP-EbC-EbP-TeC-TeP), TA-ES seems to be able to switch off a larger number of network elements (5-10%) guaranteeing the same, or slightly higher, congestion levels.

The results obtained with GA-ES and TA-ES on the Telstra network (Te10-Te20-Te30-Te40-Te50-TeC-TeP) are almost the same. This is probably due to the peculiar characteristics of the Telstra network, where the nodes average degree is very low (there is a large number of leafs, i.e., nodes with only one incident link).

Looking at Table III, it is also very interesting to observe that, in case of low workload conditions, GA-ES and TA-ES provide feasible solutions that are really close to the bounds. As bounds on the number of links and node switched-off, we use the sub-networks derived (at the optimum) by the switching off stage of TA-ES at the first iteration. This is the minimal sub-network able to satisfy all the traffic demands with fully splittable routing (that is obviously more efficient than shortest path routing). The difference in efficiency between a fully splittable routing and a shortest path routing is expected to increase for higher traffic levels. This is confirmed by the larger difference between solutions and bounds, highlighted by the computational results concerning the moderate workload conditions instances (Ex30-Ex40-Ex50-Eb30-Eb40-Eb50-ExC-ExP-EbC-EbP-TeC-TeP).

Data reported in [20] indicate that the networks are actually underutilized also during the peak workload periods. The results obtained with GA-ES and TA-ES thus suggest that it would be possible to save more than 50% of the energy also during day-hours.

C. Computing Times

Another important aspect to discuss is computing time. GA-ES and TA-ES differ drastically in terms of computational load. GA-ES is very fast (less than 5 minutes) while TA-ES is very heavy (up to 10 hours). For this reason we do not report computing times in Table III.

The above considerations indicate that it would be reasonable to use GA-ES in case of low workload conditions (finds extremely fast solutions of the same quality as TA-ES), and TA-ES in case of moderate or high workload conditions (finds solutions of better quality than GA-ES, but much slower). As far as scalability is concerned, the computing times indicate that GA-ES is more scalable, and therefore more suitable to large-size instances.

				GA-ES			TA-ES			BOUND	
Inst	Core$_{nod}$	Edge$_{nod}$	Links	N_{off}	L_{off}	Cong	N_{off}	L_{off}	Cong	N_{off}	L_{off}
Ex10	41	38	294	37 (90%)	217 (74%)	207353	37 (90%)	218 (74%)	246280	37 (90%)	218 (74%)
Ex20	41	38	294	35 (85%)	193 (66%)	269770	37 (90%)	200 (68%)	269283	37 (90%)	202 (69%)
Ex30	41	38	294	31 (76%)	169 (57%)	381939	33 (80%)	181 (62%)	517786	36 (88%)	192 (65%)
Ex40	41	38	294	25 (61%)	139 (47%)	470787	27 (66%)	149 (51%)	574416	31 (76%)	166 (56%)
Ex50	41	38	294	18 (44%)	108 (37%)	618591	23 (56%)	119 (40%)	699362	27 (66%)	149 (51%)
ExC	41	38	294	30 (73%)	184 (63%)	382263	34 (83%)	195 (66%)	407696	35 (85%)	205 (70%)
ExP	41	38	294	26 (63%)	174 (59%)	367575	32 (78%)	186 (63%)	414098	34 (83%)	200 (68%)
Eb10	56	31	322	37 (66%)	226 (70%)	105815	37 (66%)	225 (70%)	94397	37 (66%)	226 (70%)
Eb20	56	31	322	36 (64%)	219 (68%)	156496	36 (64%)	220 (68%)	241214	36 (64%)	220 (68%)
Eb30	56	31	322	33 (59%)	207 (64%)	290129	33 (59%)	204 (63%)	290129	33 (59%)	209 (65%)
Eb40	56	31	322	28 (50%)	184 (57%)	359779	30 (54%)	190 (59%)	397715	32 (57%)	201 (62%)
Eb50	56	31	322	19 (34%)	159 (49%)	531337	23 (41%)	171 (53%)	519214	27 (48%)	183 (57%)
EbC	56	31	322	26 (46%)	185 (57%)	363975	31 (55%)	199 (62%)	449094	31 (55%)	205 (64%)
EbP	56	31	322	23 (41%)	178 (55%)	357366	30 (54%)	193 (60%)	448017	32 (57%)	205 (64%)
Te10	39	65	302	19 (49%)	139 (46%)	224973	19 (49%)	139 (46%)	209092	19 (49%)	139 (46%)
Te20	39	65	302	19 (49%)	139 (46%)	858057	19 (49%)	139 (46%)	680837	19 (49%)	139 (46%)
Te30	39	65	302	17 (44%)	129 (43%)	987559	17 (44%)	127 (42%)	899394	17 (44%)	129 (43%)
Te40	39	65	302	16 (41%)	118 (39%)	1313596	16 (41%)	122 (40%)	1354361	16 (41%)	125(41%)
Te50	39	65	302	12 (31%)	110 (36%)	1791535	12 (31%)	109 (36%)	1666489	13 (33%)	113(37%)
TeC	39	65	302	16 (41%)	130 (43%)	524826	17 (44%)	129 (43%)	509236	17 (44%)	132 (44%)
TeP	39	65	302	16 (41%)	129 (43%)	534387	17 (44%)	129 (43%)	487979	17 (44%)	132 (44%)

Table III: Computational results. Ex, Eb, and Te are the abbreviations for, respectively, Exodus, Ebone and Telstra. Letters C and P, and the numbers (10-20-30-40-50) after the networks acronyms correspond to the traffic matrix considered, C indicates a constant matrix, P a poisson matricx, while the number indicates the scaling parameter used for the reduction of the maximum supported LP matrices (50 means that the maximum LP matrix has been multiplied for 0.5). For GA-ES only the best results obtained using all the sorting policies are reported.

VII. CONCLUDING REMARKS

We have proposed a first approach for energy-aware management of IP networks, by focusing attention on OSPF link weights optimization. We have shown how an efficient configuration of the OSPF link weights can allow to switch off up to 80% of the core nodes during low workload periods, and up to 65% during moderate workload periods, while guaranteeing the same point-to-point service quality and a moderate increase in network congestion. We leave as future work the design of an improved algorithm combining the complementary features of our greedy and two-phase ILP-based methods. It would be also interesting to introduce in the procedures constraints for the single link failure management.

REFERENCES

[1] J.G. Koomey. Estimating total power consumption by servers in the US and the world, 2007.

[2] GreenTouch consortium.

[3] M. Gupta and S. Singh. Greening the internet. In *ACM SIGCOM '03*, 2003.

[4] H. Mellah and B. Sansò. Review of facts, data and proposals for a greener Internet. In *Proceedings of Broadnets09*, September 2009.

[5] M. MINAMI and H. Morikawa. Some open challenges for improving the energy efficiency of the internet. In *Proc. 3rd International Conference on Future Internet (CFI 2008)*, 2008.

[6] R. Bolla, R. Bruschi, F. Davoli, and A. Ranieri. Performance constrained power consumption optimization in distributed network equipment. In *1st Workshop on Green Communications, GreenCom '09*, 2009.

[7] B. Fortz and M. Thorup. Internet traffic engineering by optimizing OSPF weights. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 519–528. IEEE, 2002.

[8] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen. Reducing the energy consumption of Ethernet with adaptive link rate (ALR). *IEEE Transactions on Computers*, pages 448–461, 2007.

[9] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall. Reducing network energy consumption via sleeping and rate-adaptation. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, pages 323–336. USENIX Association, 2008.

[10] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright. Power awareness in network design and routing. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 457–465. IEEE, 2008.

[11] L. Chiaraviglio, M. Mellia, and F. Neri. Reducing power consumption in backbone networks. In *Communications, 2009. ICC'09. IEEE International Conference*, pages 1–6. IEEE, 2009.

[12] N. Vasić and D. Kostić. Energy-aware traffic engineering. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*, pages 169–178. ACM, 2010.

[13] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini. An Energy Saving Routing Algorithm for a Green OSPF Protocol. In *IEEE INFOCOM 2010*, 2010.

[14] A. Sridharan, R. Guerin, and C. Diot. Achieving near-optimal traffic engineering solutions for current ospf/is-is networks. *Networking, IEEE/ACM Transactions on*, 13(2):234–247, 2005.

[15] M. Pióro and D. Medhi. *Routing, flow, and capacity design in communication and computer networks*. Morgan Kaufmann Publishers, 2004.

[16] A. Tomaszewski, M. Pióro, M. Dzida, M. Mycek, and M. Zagozdzon. Valid inequalities for a shortest-path routing optimization problem. In *3rd International Network Optimization Conference. Université Libre de Bruxels, Brussels*, 2007.

[17] I. Ghamlouche, T.G. Crainic, and M. Gendreau. Cycle-based neighbourhoods for fixed-charge capacitated multicommodity network design. *Operations Research*, 51(4):655–667, 2003.

[18] Rocketfuel Project [Online]. www.cs.washington.edu/research/networking/rocketfuel.

[19] Totem Project [Online]. http://totem.run.montefiore.ulg.ac.be.

[20] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and SC Diot. Packet-level traffic measurements from the Sprint IP backbone. *Network, IEEE*, 17(6):6–16, 2003.