Energy-aware traffic engineering with elastic demands and MMF bandwidth allocation

Edoardo Amaldi, Antonio Capone
Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano
Milano, Italy
amaldi@elet.polimi.it
capone@elet.polimi.it

Stefano Coniglio
Lehrstuhl II für Mathematik
RWTH Aachen University
Aachen, Germany
coniglio@math2.rwth-aachen.de

Luca G. Gianoli
Dept. of Electrical Engineering
École Polytechnique de Montréal
Montréal, Canada
luca-giovanni.gianoli@polymtl.ca

***Abstract*—In recent years, there has been a remarkable growth of the Internet energy consumption, which is expected to persist in the future at an even higher pace. At the same time the network access capacity of individual subscribers is rapidly reaching values high enough to move the traffic bottleneck from the access network to the core network in most scenarios. This will soon make the elastic nature of traffic an important aspect of network resource management and will require a redesign of the energy-aware traffic engineering techniques so far based on inelastic traffic demands.**

We propose a novel optimization approach to select a routing path for each elastic traffic demand and decide which routers and links to put to sleep so as to maximize a network utility measure depending on the traffic demand rates, while satisfying a constraint on the total energy consumption. Bandwidth is allocated to each elastic demand according to the Max-Min Fairness (MMF) paradigm, which approximates the resource allocation of the transport layer.

I. Introduction

In recent years, great effort has been devoted to increasing the access rate of Internet users, that, through the deployment of Next-Generation Optical Access networks (NGOA), will soon reach very high values (up to 2.4 Gbps in both downstream and upstream directions [1], with peak and average per-user rates of around 1 and 0.3 Gbps [8]). The access rate increase will lead to significant differences on how elastic traffic, namely all traffic transported by protocols like TCP that adapt transmission rate according to available capacity, affects the resource management in the networks of the Internet Service Providers (ISP). If the available capacity is no longer limited by the access network, connection rates will be adapted based on the backbone link capacities and how they are shared. Generally speaking, elastic demands (which, nowadays, appear to be inelastic with an upper rate limited by the access capacity), will actually compete for the backbone resources, with a variable rate depending on the routing.

Among the network and resource management issues, stretegies to reduce the energy consumption of the network infrastructure are receiving an increasing attention by device manufacturers and ISPs [13]. According to recent reports, the energy consumption of the ICT sector was of 156 GW in 2007, with 22 GW consumed exclusively by network equipments (excluding servers in data centers), and an estimated annual growth rate of 12% [21]. Other data on the consumption of some major Internet Service Providers (ISP) show for 2009 yearly consumptions up to 11 GWh. Several techniques for energy-aware network management have been recently proposed [7]. However, all these approaches neglect the impact of elastic traffic on resource sharing and assume predictable traffic demands.

In this paper, we focus on energy-aware network management strategies for backbone IP networks with elastic traffic demands and MMF flow allocation. We propose a novel approach to optimize the network routing so as to put into a low-consumption sleeping state the unnecessary routers and links and maximize a utility function related to the bandwidth allocated to each network flow. Since elastic demands require an appropriate bandwidth allocation scheme (rates are not given a priori), we consider the so-called *max-min fair* (MMF) paradigm. In a MMF flow allocation rates are assigned to traffic flows so that there is no way to give more bandwidth to any flow without decreasing the rate of another flow receiving less or equal bandwidth [17]. As widely known, an MMF allocation is the objective of most flow-control transport protocols [15] and is approximately achieved by TCP [3], [14] when connections have similar round-trip times (RTT).

Along the line of our recent work on general traffic engineering (with no energy consumption issues) [4], [6], we consider MMF allocation as a problem constraint rather than as an optimization objective. Our approach is based on a bi-level optimization problem where, at the upper level, a routing path is assigned to each demand to maximize a utility function (traffic engineering operated by the ISP), and, at the lower one, the rate of each flow is adjusted according to the MMF paradigm (resource allocation made by the transport protocol).

As to energy related issues, the choice of a sleep-based strategy where links and routers can be put to sleep is motivated by the fact that the consumption of current network devices, once activated, is almost independent of the load levels [10]. As to the routing scheme, we consider a per-flow single-path routing, like that adopted in the very popular MPLS protocol, which avoids packet reordering issues and allows for a natural definition of MMF allocation.

The paper is organized as follows. After summarizing related work in Section II, in Section III we describe our approach and present a mixed-integer linear programming (MILP) formulation. In Section IV, we propose a restricted-path heuristic algorithm and, in Section V, we report some computational results. Finally, Section VI contains some concluding remarks.

II. Related work

Although a large body of work on energy-aware network management has appeared in the recent years [7], we are not aware of articles that jointly consider energy consumption issues and elastic traffic demands with MMF allocation.

In the literature on energy-aware network management with inelastic traffic demands, great attentions has been devoted to the development of traffic engineering strategies to save energy by putting to sleep routers and links [2], [5], [9], [11] or only links [12], [23]. We mention greedy heuristics to switch-off network network elements in flow-based routing networks [11], exact and heuristic methods to optimize the daily network consumption in MPLS networks [2], MILP-based algorithms to achieve energy savings by efficiently configuring the link weights of the Open Shortest Path First (OSPF) protocol [5], a modified version of OSPF that reduces the number of active links by exploiting shared shortest path trees [12], a MILP formulation based on pre-computed paths for networks with both MPLS and OSPF [23], and approaches for the energy-aware optimization of routing trees used by the Carrier Grade Ethernet protocol [9].

To the best of our knowledge, in all previous work related to MMF allocation of elastic demands, the MMF criterion is considered as the objective function, and no energy consumption issues are taken into account. When the routing path of each elastic demand is already given, a well-known algorithm (the so-called *Progressive filling* method) can be used to find an MMF allocation vector [17] in polynomial-time. For the case where both the paths and the MMF allocation have to be determined, techniques have been proposed to determine a network routing which admits an MMF flow allocation assuming both unsplittable [17], [20] or splittable [16], [17] routing. In general, given a nominal optimization problem, a solution to its version where we look for an MMF allocation can be found by solving a sequence of problem which are easily derived from the original one [19]. Unlike in the above mentioned work, in [4], [6] we were the first to propose to consider the MMF allocation of elastic demands as constraint in a more general traffic engineering problem aiming at maximizing a utility function selected by the network operator (e.g. the maximization of the total network throughput).

III. Problem description and optimization model

In this work we propose and investigate the following energy-aware traffic engineering problem subject to MMF flow allocation. Given a network with link capacities and a set of elastic communications specified by the corresponding origin-destination pairs, decide which routers and links can be put to sleep mode and how to route the elastic traffic demands so as to maximize a utility function (depending on the bandwidth allocated to all the communications) subject to an energy budget constraint, while assuming that each demand is routed along a single path and that the bandwidth is allocated according to the MMF principle.

As it is well known, the per-flow scheme with single path (unsplittable) routing is typically adopted by the MPLS routing protocol. Although a deviation between TCP and MMF may occur in the presence of communications with different round-trip-times (RTT) [14], other congestion avoidance mechanisms such as *random early detection* (RED), commonly implemented in IP routers, allow to effectively limit this phenomenon [3].

Before describing a MILP formulation for the above problem, we need to briefly recall the definition and characterization of MMF bandwidth allocation. When a single routing path is selected for each elastic demand, a bandwidth (flow) allocation is MMF if there is no way to assign more bandwidth to any communication without decreasing the rate of another communication receiving less or equal bandwidth. Equivalently, an allocation vector is MMF if each elastic demand d is routed through at least a *bottleneck link* (i, j), i.e., a link (i, j) which is saturated (with no residual capacity) and such that the flow of all the other elastic demands through that link does not exceed the flow of demand d.

Since for each set of unsplittable routing paths there exists a single MMF allocation vector, once a set of paths is given (i.e., once it has been selected by the ISP network operator) the rate of each communication is uniquely determined by the transport protocol (TCP) which allocates bandwidth according to the MMF principle.

A. MILP formulation

Let $G = (V, A)$ be a directed graph representing an IP network topology where V and A correspond to the set, respectively, of routers and links. Let also $V_s := V \setminus \{s\}$. Let $c_{ij} \geq 0$ denote the capacity of link $(i, j) \in A$, and K be the set of elastic demands, where $(s, t) \in K$ is the demand between node s and node t. Furthermore, let the positive continuous variable ϕ^{st} denote the data rate of communication $(s, t) \in K$, and Δ_{st} be the number of different connections of which each traffic demand (s, t) is composed. This parameter allows to generalize the model in [?] to the case where multiple connections have the same origin-destination pairs.

As utility function, we consider either the total weighted throughput

$$\max \sum_{(s,t)\in K} w_{st}\phi^{st} \quad (1)$$

with real weights w_{st}, or

$$\max \sum_{(s,t)\in K} w_{st}\Delta_{st}\alpha(1-e^{\frac{1}{\beta}\frac{\phi^{st}}{\Delta_{st}}}), \quad (2)$$

for suitable $\alpha, \beta > 0$, which allows to split more equally the bandwidth between the different demands since it introduces a saturating effect for large ϕ^{st}. By adding a continuous variable per demand and a linear constraint per piece, we obtain a linear programming piecewise-affine approximation (with 6 pieces) of the above concave functions, see e.g. [22].

Besides the flow conservation and capacity constraints

$$\sum_{(i,j)\in A} f_{ij}^{st} - \sum_{(j,i)\in A} f_{ji}^{st} = \begin{cases} \phi^{st} & \text{if } i=s \\ -\phi_{st} & \text{if } i=t \\ 0 & \text{else} \end{cases} \quad \forall i \in V,\ (s,t)\in K \quad (3)$$

$$\sum_{(s,t)\in K} f_{ij}^{st} \le c_{ij} \quad \forall (i,j)\in A,\ (s,t)\in K \quad (4)$$

$$f_{ij}^{st} \le c_{ij}x_{ij}^{st} \quad \forall (i,j)\in A,\ (s,t)\in K, \quad (5)$$

where f_{ij}^{st} is a positive continuous variable representing the amount of flow of demand $(s,t)\in K$ on link $(i,j)\in A$ and x_{ij}^{st} is a binary variable which is equal to 1 if $f_{ij}^{st} > 0$, we introduce the following groups of constraints to correctly define a single routing path for each demand:

$$\sum_{(i,j)\in A} x_{ij}^{st} \le \begin{cases} 0 & \text{if } i=t \\ 0 & \text{if } j=s \\ 1 & \text{else} \end{cases} \quad \forall i \in V,\ (s,t)\in K \quad (6)$$

$$u_{ijh}^{st} \le x_{ij}^{st} \quad \forall h \in V_s,\ (s,t)\in K,\ (i,j)\in A \quad (7)$$

$$\sum_{(i,h)\in A} x_{ih}^{st} = z_h^{st} \quad \forall h\in V_s,\ (s,t)\in K \quad (8)$$

$$\sum_{(i,j)\in A} u_{ijh}^{st} - \sum_{(j,i)\in A} u_{jih}^{st} = \begin{cases} z_h^{st} & \text{if } i=s \\ -z_h^{st} & \text{if } i=t \\ 0 & \text{else} \end{cases} \quad \forall h\in V_s, j\in V, (s,t)\in K, \quad (9)$$

where z_h^{st} and u_{ijh}^{st} are binary variables equal to 1 if, respectively, node h is on the path used to satisfy demand (s,t) and if link (i,j) belongs to the subpath connecting nodes s and h which is derived from the used by demands (s,t). By sending an auxiliary unit of flow from the origin of each demand to every intermediate node of the corresponding routing path, Constraints (6)-(9) allow to prevent subtours.

A further set of constraints is used to evaluate the power-state of both routers and links:

$$x_{ij}^{st} \le l_{ij}, \quad \forall (i,j)\in A,\ (s,t)\in K \quad (10)$$

$$l_{ij} \le \xi_i, \quad \forall (i,j)\in A \quad (11)$$

$$l_{ij} \le \xi_j, \quad \forall (i,j)\in A \quad (12)$$

$$\sum_{i\in V} \varrho_i\xi_i + \sum_{(i,j)\in A} p_{ij}l_{ij} \le B, \quad (13)$$

where l_{ij} and ξ_i are binary variables equal to 1 if, respectively, link (i,j) or node i are fully activated. The positive parameters p_{ij} and ϱ_i represent the power consumption of link (i,j) and node i. Constraints (10) forbid to route a demand through a sleeping link, Constraints (11)-(12) force to put into sleeping status a link connected to sleeping nodes, and Constraint (13) imposes a maximum total energy consumption of B.

To guarantee an MMF allocation, we add the following constraints, which are a restated version of those in [20]:

$$\sum_{(i,j)\in A} y_{ij}^{st} \ge 1 \quad \forall (s,t)\in K \quad (14)$$

$$\sum_{(o,d)\in K} f_{ij}^{od} \ge c_{ij}y_{ij}^{st} \quad \forall (i,j)\in A,\ (s,t)\in K \quad (15)$$

$$\vartheta_{ij} \ge f_{ij}^{st}/\Delta_{st} \quad \forall (i,j)\in A,\ (s,t)\in K \quad (16)$$

$$f_{ij}^{st}/\Delta_{st} \ge \vartheta_{ij} - c_{ij}(1-y_{ij}^{st}) \quad \forall (i,j)\in A,\ (s,t)\in K, \quad (17)$$

where y_{ij}^{st} is the binary variable equal to 1 if link (i,j) is a bottleneck for the commodity (s,t) and ϑ_{ij} is the positive continuous variable used to identify the largest flow on link (i,j) (Constraints (16)). Constraints (14) ensure that we have at least a bottleneck arc for each (s,t) pair, while Constraints (15) guarantee that the bottleneck arcs are saturated. Finally, Constraints (16)-(17) impose that the flow through a bottleneck arc (i,j) for a pair (s,t) be at least as large as the flow through (i,j) for all the other origin-destination pairs.

To accelerate the convergence of the MILP solver to an optimal solution, we also include the following two groups of simple valid inequalities which yield tighter LP relaxations:

$$y_{ij}^{st} \le x_{ij}^{st} \quad \forall (i,j)\in A,\ (s,t)\in K \quad (18)$$

and

$$\frac{\phi^{st}}{\Delta_{st}} \ge \frac{\min_{(i,j)\in A}\{c_{ij}\}}{\sum_{(o,d)\in K}\Delta_{od}} \quad \forall (s,t)\in K. \quad (19)$$

Constraints (19) are valid since any MMF flow for a demand (s,t) saturates at least a link and, due to the MMF allocation, the smallest possible bandwidth arises when the link is shared by all the communications.

IV. The restricted-path heuristic

Since the above MILP formulation turns out to be very challenging to solve especially when the number of demands increases, we propose a heuristic based on a restricted-path formulation which focuses on a subset of heuristically pre-computed paths. To describe such formulation, let P^{st} denote the set of pre-computed paths for the demand (s,t) and let σ_p^{st} be a binary variable equal to 1 if the path $p \in P^{st}$ is used to satisfy the demand (s,t), and 0 otherwise. To guarantee that a single path is used for each demand, we add:

$$\sum_{p \in P^{st}} \sigma_p^{st} = 1, \quad \forall (s,t) \in K. \tag{20}$$

Furthermore, compared to (3)-(19), Constraints (6)-(9) are discarded, while Constraints (5), (10), and (18) are modified as follows to take into account the pre-computed paths:

$$f_{ij}^{st} \le c_{ij} \sum_{p \in P^{st}:(i,j) \in p} \sigma_p^{st} \quad \forall (i,j) \in A, \, (s,t) \in K \tag{21}$$

$$\sigma_p^{st} \le l_{ij}, \quad \forall (i,j) \in A, \, (s,t) \in K, \; p \in P^{st} : (i,j) \in p \tag{22}$$

$$y_{ij}^{st} \le \sum_{p \in P^{st}:(i,j) \in p} \sigma_p^{st} \quad \forall (i,j) \in A, (s,t) \in K. \tag{23}$$

The pre-computed paths are generated with the following procedure. For each demand $(s,t) \in K$, we first compute with an LP formulation the maximum flow m_{st} that can be routed from node s to t with unit capacity on the links. m_{st} is equal to the the maximum number of link-disjoint paths between s and t. The path computation is then run in three different steps. A first set of paths (one path for each demand) is extracted from the minimal energy directed Steiner-tree (computed with a MILP formulation) that connects all the nodes that are source or destination of traffic. A second set is computed with an iterative procedure (ω iterations, for some $\omega \in \mathbb{N}$), where, at each iteration, we first assign a random weight to each link, and then compute m_{st} link-disjoint shortest paths for each demand (s,t). At the end of the routine, we obtain ωm_{st} shortest paths for each demand. The last group of paths is computed by generating ω sets of random link weights, which we use to obtain ω different minimum cost Steiner-trees. A single path for each demand is then extracted from each Steiner-tree (ω paths for each demand). Thus, we obtain $1 + \omega m_{st} + \omega$ paths for each commodity. Note that the motivation for our two path generation strategies is to find both disjoint paths (which allow to better distribute the traffic so as to maximize the objective function) as well as paths which are highly correlated, i.e., sharing many links (so as to minimize the energy consumption).

V. Computational results

Computational experiments have been carried out on four network topologies provided by the SND library [18]: `polska`, `abilene`, `geant` and `nobel-eu`. Network routers are equipped with M10i chassis, while each link is provided with a line card randomly chosen from those of Table I.

Table I
Router Chassis and Link Cards

case	device	capacity	hourly power cons.
all	Chassis Juniper M10i	16Gbps	86.4 W
1	Gigabit-Eth 1 port	2 Gbps	7.3 W
2	Fast-Eth 12 ports	2.4 Gbps	18.6 W
3	Gigabit-Eth 4 ports	8 Gbps	31 W
4	SONET/SDH OC-48c	5 Gbps	41.4 W

Table II
Instances testbed

Instances			**1**		**2**		**3**		**4**		**5**	
Net	$\lvert V\rvert$	$\lvert A\rvert$	$\lvert V_e\rvert$	$\lvert K\rvert$	$\lvert V_e\rvert$	$\lvert K\rvert$	$\lvert V_e\rvert$	$\lvert K\rvert$	$\lvert V_e\rvert$	$\lvert K\rvert$	$\lvert V_e\rvert$	$\lvert K\rvert$
pol	12	36	4	6	5	10	6	21	7	28	8	36
abi	15	44	4	12	5	20	6	30	7	42	8	56
gea	22	72	5	20	6	30	7	42	8	56	9	72
nob	28	82	9	36	10	45	11	55	12	66	13	78

The number of chassis in each node varies according to the total capacity and the number of ports required by the cards connected to it. Two different random equipment configurations have been tested for each topology. Furthermore, for each network topology we consider a varying number of randomly selected *edge nodes*, i.e., nodes that cannot be put to sleep because they are source or destination of at least one traffic demand. As for the set of commodities, we consider the traffic matrices provided by the SNDlib and discard all the demands generated between non-edge nodes. For each demand set, we generate two different scenarios where Δ_{st} is randomly chosen in $[0, 10]$ for each $(s,t) \in K$. Network instances are summarized in Table II, where $|V|$, $|A|$, $|V_e|$, and $|K|$ represent, respectively, the number of nodes, unidirectional links, edge nodes, and elastic demands. Our MILP formulations are solved with CPLEX 12.4.0.1 using the AMPL modeling language. Computational experiments are conducted on 4 Intel i7 processors and 8 GB of RAM. A time limit of 1 hour is imposed. For simplicity, we assume that all the communications have the same weight $w^{st} = 1$. In the objective function (2), we let $\alpha = 1000$ and $\beta = 200$.

First, we experiment with the MILP formulation (3)-(19) and the restricted-path heuristic on the smallest networks, namely, `polska` and `abilene`, considering both objectives (1) and (2). The results obtained with the MILP formulation are reported in Figure 1 (subfig. (a)-(b)-(c)-(d)). The four plots show the relationship between network utility (Y-axis) and energy consumption (X-axis). For each network instance (*network name-number of edges*) we considered five different energy budget values B, from the minimum energy consumption needed to connect all the edge nodes up to the consumption of the full active network. Both energy and utility values are normalized w.r.t. to those computed for the solutions obtained with the minimum B. Furthermore, the values are averaged over the four values registered when experimenting with the four combinations obtained by mixing the two configurations and the two traffic scenarios generated for each network instance. As expected, the increment of the energy budget produces a natural increase of the network utility. With obj. (1), the utility increases up to 2.5 and 2.8 times w.r.t. the solutions computed by imposing the minimum consumption budget in `polska` and `abilene`, respectively. The increase is not as large with obj. (2), with an increment up to 1.8 and 1.65 times, arguably because of the concavity of (2). Note that the X-extension reduction observed in the curves representing the instances with an higher number of edge nodes is due to the smaller consumption difference between the minimum topology (that with the smallest B) and the full active one. In fact, with more edge-nodes, a fewer routers can be put to sleep. Finally, the last point of each curve represents the network utility obtained with a *standard configuration*, according to which all the traffic demands are routed in full active network with a shortest path scheme and all the link weights set equal to $1/c_{ij}$ (the weight setting typically used by the operators). Observe that the network utility obtained by the standard configurations can be achieved by the optimized solutions also when imposing a reduced energy budget. Furthermore, we note an increase between 10% and 30% of the utility for the same power consumption (both fully activated networks) when the optimization is performed. In Table III, we report other interesting results concerning both the exact and the heuristic methods (with $\omega = 4$). Columns g_o, g_h, t_m and t_h represent, respectively, the gap of the solution obtained with the complete formulation w.r.t. the best upper bound, the gap between the heuristic solution and that of the formulation, the computational times needed to solve the complete formulation and that for the restricted-path algorithm. Note that, on average, the value for g_o is greater than 0 because CPLEX is run with a time-limit (tl) of 3600 s and, not always, optimal solutions are found within such time limit. With obj. (1), the average and 90-percentile g_o are respectively around 2.5% and 7% with both `polska` and `abilene`. A significant improvement is obtained with obj. (2), for which the 90-percentile value for `polska` is as small as 0.4%. This behaviour is likely due to the saturation effect induced by obj. (2), which helps the resolution algorithm to move towards an MMF allocation. As for the restricted-path heuristic, we register a 90-percentile g_h around 6% with (1), and lower than 3.2% with (2). This result confirms the validity of the heuristic approach, showing that even by using a restricted set of paths we can obtain close to optimal solutions within a computing time that is, on average, 1000 seconds smaller than that required when solving the complete formulation.

Since, on the largest instances with more than 20 nodes, it is hard to find feasible solutions of good quality with our MILP formulation, in a second set of tests we experiment with the restricted-path heuristic on the largest `geant` and `nobel-eu` networks, with ω equal to $2, 4$ and a time-limit of one hour. Table IV summarizes the information that allows to evaluate the heuristic performance and the impact of the number of pre-computed paths given as input

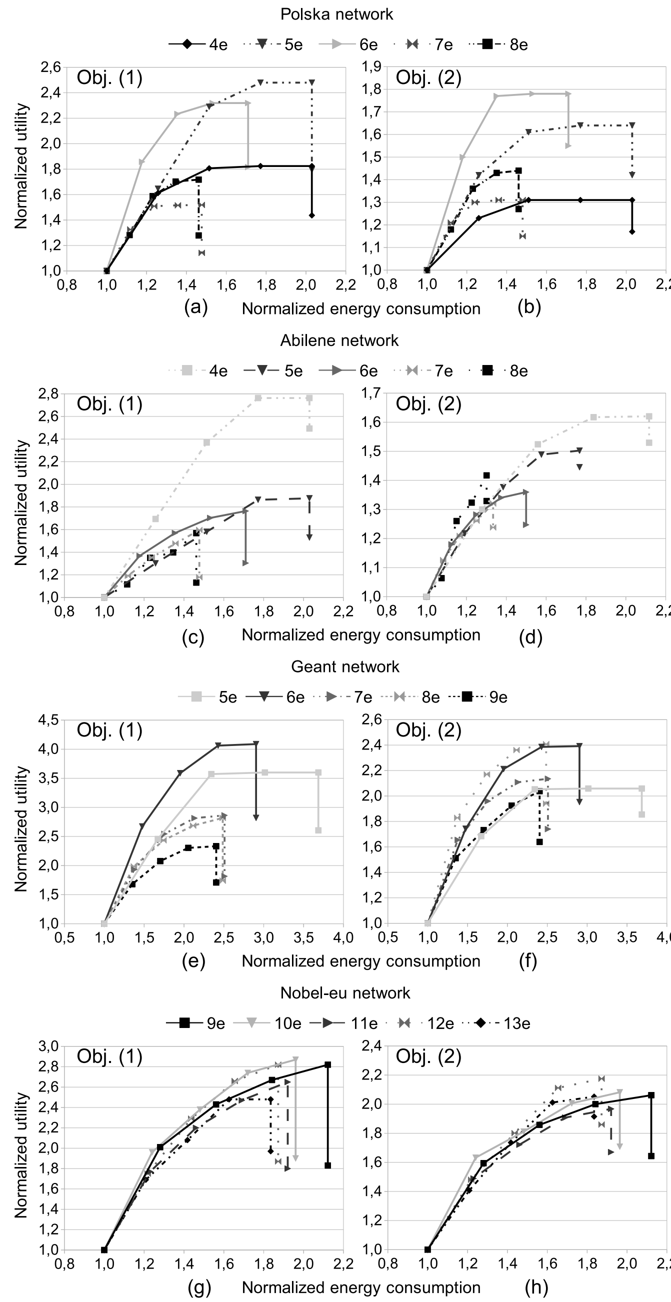

Figure 1. Computational results obtained with the MILP formulation for `polska` and `abilene`, and with the restricted-path heuristic for `geant` and `nobel-eu`.

to the procedure. Columns $g_{\omega 2}^{\omega 4}$, $g_{sp}^{\omega 2}$, and $g_{sp}^{\omega 4}$ represent, respectively, the gap between the heuristic with $\omega = 2$ and $\omega = 4$, the gap between the network utility obtained with a classic shortest path routing and that for the restricted-path heuristic with $\omega = 2$ and $\omega = 4$ when the network is fully activated. In `geant`, doubling the number of paths produces an average utility increase of 2.5% and 1.3% with obj. (1) and (2), respectively. A slight reduction, around 0.5% on average, is instead observed in `nobel-eu`. Since the addition of the new path variables tends to increase the problem complexity, higher time-limits may

Table III
Comparison between the solutions found by the MILP formulation and the restricted-path heuristic: time-limit of 1 hour, $\omega = 4$.

	polska				**abilene**			
obj-(1)	g_o	g_h	t_m	t_h	g_o	g_h	t_m	t_h
mean	2.8%	1.8%	2.5k	0.3k	1.6%	2.4%	1.3k	0.7k
st.dev	4.9 %	3.9 %	2.8k	0.3k	2.9%	3.2%	1.6k	1.2k
90perc	6.6%	5.3%	tl	0.8k	6.3%	7.3%	tl	tl
obj-(2)	g_o	g_h	t_m	t_h	g_o	g_h	t_m	t_h
mean	0.2%	0.7%	1.2k	70	1.0%	1.0%	1.3k	0.8k
st.dev	0.6 %	1.3 %	1.7k	0.4k	2.1%	2.7%	1.6k	1.4k
90perc	0.4%	3.2%	tl	0.2k	3.9%	3.1%	tl	tl

Table IV
Gap between the network utility of the solutions found by the restricted-path heuristic with $\omega = 2, \omega = 4$, and standard shortest path routing when the network is fully activated.

	geant			**nobel-eu**		
obj-(1)	$g^{\omega 4}_{\omega 2}$	$g^{\omega 2}_{sp}$	$g^{\omega 4}_{sp}$	$g^{\omega 4}_{\omega 2}$	$g^{\omega 2}_{sp}$	$g^{\omega 4}_{sp}$
mean	-2.5%	30.9%	33.4%	0.4%	31.0%	30.8%
st.dev	2.8%	5.0 %	5.4%	5.5%	9.5%	10.0%
obj-(2)	$g^{\omega 4}_{\omega 2}$	$g^{\omega 2}_{sp}$	$g^{\omega 4}_{sp}$	$g^{\omega 4}_{\omega 2}$	$g^{\omega 2}_{sp}$	$g^{\omega 4}_{sp}$
mean	-1.3%	17.4%	18.7%	0.6%	15.9%	14.7%
st.dev	2.8%	7.4%	7.3%	3.6%	7.4%	9.2%

be required to reach solutions of comparable quality. Note that, on average, a utility improvement of around 33% and 18% is observed with obj. (1) and (2), respectively, when comparing the optimized solutions and the standard ones with the full active networks. To conclude, in Figure 1 (subfig. (e)-(f)-(g)-(h)) we report the budget-utility trend for the solutions obtained with $\omega = 4$. The same behaviour is observed for the largest instances.

VI. Concluding remarks

We have proposed a novel approach to optimize network utility in IP backbone networks with elastic traffic demands subject to MMF bandwidth allocation and an energy consumption constraint. We have described an exact MILP formulation and a heuristic algorithm based on a restricted set of pre-computed routing paths. With the MILP formulation, we are able to solve to optimality instances with up to 15 nodes, 40 links and 40 demands, while with the restricted-path heuristic we obtain sub-optimal solutions for networks with up to 30 nodes, 100 links and 100 demands. The utility values of the resulting solutions are close to those obtained with a typical shortest path routing and a full active network, in spite of the fact that we impose a 35% reduction on the total energy consumption. We leave as future work the development of more efficient heuristics to tackle larger instances.

References

[1] ITU-T G.984.1 Gigabit-capable passive optical networks (GPON): General characteristics. http://www.itu.int/rec/T-REC-G.984.1-200803-I/en, March 2008.

[2] B. Addis, A. Capone, G. Carello, L.G. Gianoli, and B. Sansò. Energy management through optimized routing and device powering for greener communication networks. *Networking, IEEE/ACM Transactions on*, PP(99):1–1, 2013.

[3] E. Altman, C. Barakat, E. Laborde, P. Brown, and D. Collange. Fairness analysis of tcp/ip. In *Decision and Control, 2000. Proceedings of the 39th IEEE Conference on*, volume 1, pages 61–66. IEEE, 2000.

[4] E. Amaldi, A. Capone, S. Coniglio, and L.G. Gianoli. Network optimization problems subject to max-min fair flow allocation. *Communications Letters, IEEE*, 17(7):1463–1466, 2013.

[5] E. Amaldi, A. Capone, and L.G. Gianoli. Energy-aware ip traffic engineering with shortest path routing. *Computer Networks*, 57(6):1503 – 1517, 2013.

[6] E. Amaldi, S. Coniglio, L.G. Gianoli, and C.U. Ileri. On single-path network routing subject to max-min fair flow allocation. *Electronic Notes in Discrete Mathematics*, 41:543–550, 2013.

[7] R. Bolla, R. Bruschi, F. Davoli, and F. Cucchietti. Energy efficiency in the future internet: a survey of existing approaches and trends in energy-aware fixed network infrastructures. *Communications Surveys & Tutorials*, (99):1–22.

[8] D. Breuer, F. Geilhardt, R. Hulsermann, M. Kind, C. Lange, T. Monath, and E. Weis. Opportunities for next-generation optical access. *Communications Magazine, IEEE*, 49(2):16–24, 2011.

[9] A. Capone, D. Corti, L. Gianoli, and B. Sansó. An optimization framework for the energy management of carrier ethernet networks with multiple spanning trees. *Computer Networks*, 56(17):3666 – 3681, 2012.

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

[11] L. Chiaraviglio, M. Mellia, and F. Neri. Minimizing isp network energy cost: Formulation and solutions. *IEEE/ACM Transactions on Networking (TON)*, 20(2):463–476, 2012.

[12] A. Cianfrani, V. Eramo, M. Listanti, M. Polverini, and A. Vasilakos. An ospf-integrated routing strategy for qos-aware energy saving in ip backbone networks. *Network and Service Management, IEEE Transactions on*, PP(99):1–14, 2012.

[13] GreenTouch consortium. www.greentouch.org.

[14] J.Y. Le Boudec. Rate adaptation, congestion control and fairness: A tutorial. http://moodle.epfl.ch/file.php/523/CC_Tutorial/cc.pdf, 2005.

[15] L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. *Networking, IEEE/ACM Transactions on*, 10(3):320–328, 2002.

[16] D. Nace. A linear programming based approach for computing optimal fair splittable routing. In *Computers and Communications, 2002. Proceedings. ISCC 2002. Seventh International Symposium on*, pages 468–474. IEEE, 2002.

[17] D. Nace and M. Pióro. Max-min fairness and its applications to routing and load-balancing in communication networks: a tutorial. *Communications Surveys & Tutorials, IEEE*, 10(4):5–17, 2008.

[18] S. Orlowski, R. Wessäly, M. Pióro, and A. Tomaszewski. SNDlib 1.0 - survivable network design library. *Networks*, 55(3):276–286, 2010.

[19] M. Pioro. Fair routing and related optimization problems. In *Advanced Computing and Communications, 2007. ADCOM 2007. International Conference on*, pages 229–235. IEEE, 2007.

[20] A. Tomaszewski. A polynomial algorithm for solving a general max-min fairness problem. *European transactions on telecommunications*, 16(3):233–240, 2005.

[21] W. Vereecken, L. Deboosere, D. Colle, B. Vermeulen, M. Pickavet, B. Dhoedt, and P. Demeester. Energy efficiency in telecommunication networks. 2008.

[22] H.P. Williams. *Model building in mathematical programming.* John Wiley and Sons, 1999.

[23] M. Zhang, C. Yi, B. Liu, and B. Zhang. Greente: Power-aware traffic engineering. In *Network Protocols (ICNP), 2010 18th IEEE International Conference on*, pages 21–30. IEEE, 2010.