Les Cahiers du GERAD ISSN: 0711–2440

Robust Energy-Aware Multi-Period Traffic Engineering with Flow-Based Routing

B. Addis, A. Capone, G. Carello,
L.G. Gianoli, B. Sansò

G–2013–101

December 2013

Robust Energy-Aware Multi-Period Traffic Engineering with Flow-Based Routing

Bernardetta Addis

Department of Informatic
Università degli Studi di Torino, Italy

addis@di.unito.it

Antonio Capone
Giuliana Carello

Dipartimento di Elettronica, Informazione e Bioingegneria
Politecnico di Milano, Italy

capone@elet.polimi.it
giuliana.carello@polimi.it

Luca G. Gianoli
Brunilde Sansò

GERAD & Department of Electrical Engineering
Polytechnique Montréal, Canada

luca-giovanni.gianoli@polymtl.ca
brunilde.sanso@polymtl.ca

December 2013

Les Cahiers du GERAD

G–2013–101

Abstract: A robust multi-period model is proposed to minimize the energy consumption of IP networks, while guaranteeing the satisfaction of uncertain traffic demands. Energy savings are achieved by putting into sleep mode cards and chassis. The study of the solution robustness shows that there is a trade-off between energy consumption and the solutions conservatism degree. The model allows this trade-off to be tuned by simply modifying a single parameter per link. The multi-period optimization is constrained by inter-period limitations necessary to guarantee network stability. Both, exact and heuristic methods are proposed. Results show that up to 60% of the energy savings can be achieved for realistic test scenarios in networks operated with flow-based routing protocols (i.e. MPLS) and with a good level of robustness to traffic variations.

1 Introduction

Due to Internet rapid expansion, it is said that the ICT sector contribution is 2% (0.8 Gt CO2) of annual global greenhouse gas (GHG) emissions [1], and that in 2007 the Internet was responsible for 5.5% of the total energy consumption in the world [2]. *Green Networking* aims at optimizing telecommunication network energy consumption by working at different levels: development of i) new energy efficient network devices, ii) new methodologies for power aware network design and iii) new energy management strategies [3]. The reader is referred to [4] for a discussion on different types of Green Networking proposals.

In this paper we focus on IP network energy-aware management and we aim at limiting the energy-wise negative effects due to bandwidth over provisioning, without reducing the QoS. In fact, although network utilization varies typically from 5% (night hours) to 50% (peak hours) [5], the network consumption remains practically constant because the energy consumed by network devices is almost independent of the traffic load [6].

A promising strategy is represented by energy-aware *Traffic Engineering* (TE), carried out assuming that unused devices can be put to sleep. Energy-aware *Traffic Engineering* is strictly influenced by i) the routing protocol considered and by ii) the accuracy of the traffic estimations.

In this paper we consider IP networks operated with Multi Protocol Label Switching (MPLS), that is, together with the Open Shortest Path First (OSPF) protocol, the most popular protocol adopted in the backbone IP networks. MPLS explicitly selects the route of each individual traffic demand, guaranteeing in this way a very flexible *TE*.

Differently from previous work on IP networks energy savings, we specifically consider here the uncertainty of traffic estimations and the robustness of the network. We propose an offline method based on predicted values of traffic and on a robust optimization approach that assumes that traffic demands vary within a given uncertainty set [7]. We show that it is possible to obtain optimal solutions that satisfy network constraints even when the traffic demands do not exactly take the nominal values. We also show that there is a trade-off between robustness and energy consumption which can be tuned by just using a single parameter per link.

The problem that we address in this paper is how to optimize network energy consumptions without affecting the network performance and the efficiency of network management mechanisms. For this purpose, we propose a multi-period optimization problem where we aim at minimizing the energy consumption of IP networks over a set of time intervals, while guaranteeing the satisfaction of the uncertain traffic demands. The value of each traffic demand can vary with a uniform distribution inside a symmetric interval centered on a nominal predicted value. A per-flow single path routing scheme is considered and energy savings are achieved by putting to sleep unused routers and links. Some inter periods constraints to limit the number of device switching on along the entire set of intervals are used to guarantee network stability and to preserve the expected device lifespan. An ILP formulation and a heuristic method based on the same ILP formulation are presented to solve the problem.

The remainder of the paper is organized as follows. In Section 2 we review previous papers on green networking and point out the novelties of our work. In Section 3.1 we present the energy management strategy proposed, the system modelling assumptions and the ILP formulation, while in Section 3.2 we present the robust variant of the previous basic formulation. In Section 4 the new robust heuristic based on mathematical programming is proposed. A set of numerical results obtained on four networks are shown and discussed in Section 5. Finally, concluding remarks are exposed in Section 6.

2 Related Work

The problem of reducing Internet energy consumption has been at first presented in the seminal work by Gupta and Singh [3]. We refer the reader to [4, 8, 9] for exhaustive surveys of the research on the topic, and for accurate taxonomies to classify the different green techniques.

In particular, as for energy-aware *TE*, a limited number of papers has been presented in the last years. They can be classified according to the routing protocol. The per-flow routing considered in this paper has been previously adopted in [10, 11]. The approach to off-line energy management proposed in [10] aims at reducing network energy consumption by switching off nodes and interfaces and it is based on a greedy algorithm that considers a single set of traffic demands. On the other hand, we model a multi-period scenario with a set of uncertain traffic demands corresponding to each different time period and jointly optimize energy management in all scenarios following a robust approach. We assume a single path routing (unsplittable flows) that can be applied to MPLS-based networks and consider limitations to the state variations of devices. Some on-line Energy-Aware Traffic Engineering (EATe) techniques to optimize links and routers power consumption are instead proposed in [11]; these on-line procedures exploit a local search scheme and are based on the assumption that the energy profiles of network devices are strongly dependant on the utilization.

Networks operated with shortest path routing protocols (e.g. OSPF) are instead treated in [12–14] but none of those papers consider either multi-period optimization, demands uncertainty, or inter-period constraints, like we do. The heuristic approaches proposed in [12, 13] are based on the idea of achieving energy savings (by switching off both links and nodes) and minimizing network congestion by efficiently optimizing the link weights. The Energy Aware Routing (EAR) algorithm presented in [14] aims at putting into sleep mode the network elements by using a modified version of the OSPF protocol where traffic demands are routed along the shortest path trees computed by only a certain subset of routers. This method focuses only on the routing protocol and does not directly consider traffic demands and network capacity limitations.

Finally, we survey some recent contributions that adopt different perspectives: methods for switching off network devices in networks operated with an hybrid routing scheme (MPLS plus OSPF) [15], procedures that turn off network links by only considering network topology features (traffic demands are ignored) [16], a distributed algorithm to determine the operating configuration of each node so as to minimize energy consumption [17], and new energy-aware protocols [18, 19].

To the best of our knowledge, except for our preliminary work [20], no other works concerning energy-aware multi-period optimization with inter-period constraints have been presented yet. W.r.t. [20], where a GRASP heuristic for the basic problem (without uncertainty) is proposed, in this paper we present a new modelling framework for managing traffic uncertainty and we propose a completely novel heuristic approach. We refer the reader respectively to [21] and [7] for general survey on multi-period network optimization and robust optimization.

3 The problem

3.1 The MILP formulation

In our problem we consider: i) an IP network represented by a graph $G(N, A)$ where each router is composed of a chassis and a set of line cards. Router chassis are represented by the set of nodes N. The set of line cards connecting router $i \in N$ and router $j \in N$ is represented by the link (i, j), with n_{ij} line cards installed on link (i, j) ($n_{ij} \geq 1$). ii) A set S of time intervals σ of duration h_σ, the sum of which corresponds to an entire day, and iii) a set of uncertain traffic demands D, one for each time interval. Each uncertain traffic demand $d \in D$ is characterized by a nominal value ρ_d, and by a real non negative parameter $r_d^\sigma \in [0, 1]$ that indicates the amount of the nominal value ρ_d that has to be satisfied during scenario σ. Since real traffic demands cannot be exactly predicted, we assume r_d^σ to be uncertain, with the possibility to take values inside the symmetric interval $[\hat{r}_d^\sigma - \bar{r}_d^\sigma, \hat{r}_d^\sigma + \bar{r}_d^\sigma]$. Note that the values of $\hat{r}_d^\sigma$ assigned to each demand follow exactly the profile shown in Figure 1. Following the idea proposed in [7], we use the parameter $\Gamma_{ij}^\sigma \in [0, |D|]$ to limit the number of traffic demands that are considered uncertain: in this way we can easily tune the degree of conservatism of the solution. From a practical point of view we use $\Gamma_{ij}^\sigma$to exclude the unlikely situations where all the traffic demands routed on a link (i, j) assume the maximal values simultaneously in scenario σ. Although the set of demands may be different for each link, by allowing the MILP solution to choose the set of demands which assume the maximal value for each link, we guarantee that the solution is feasible for the worst possible condition, while keeping the overall model manageable.

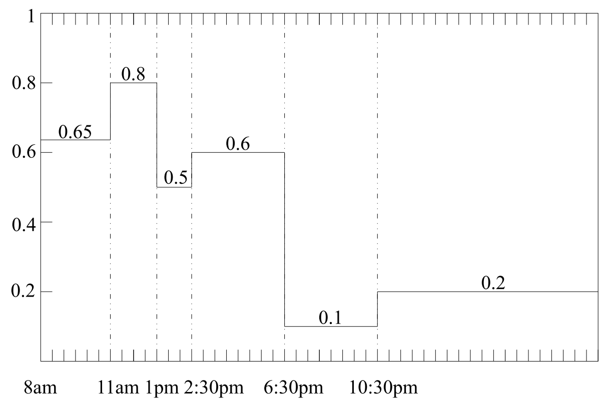

Figure 1: Traffic scenarios.

The target of our optimization is the minimization of the network energy consumption, by putting in sleep mode unnecessary line cards and chassis. The objective function can be expressed as:

$$\min \sum_{\sigma \in S} \left(\sum_{j \in N} \bar{\pi} y_j^\sigma \right) h_\sigma + \sum_{\sigma \in S} \left(\sum_{(i,j) \in A} \pi_{ij} w_{ij}^\sigma \right) h_\sigma + \sum_{\sigma \in S} \sum_{j \in N} z_j^\sigma \quad (1)$$

where w_{ij}^σ are integer variables in $\{0, \ldots, n_{ij}\}$ that represent the number of line cards activated on link (i, j) during scenario σ, y_j^σ are binary variables that are equal to 1 when chassis j is on during scenario σ, z_j^σ are non negative continuous variables, which represent the energy consumption if chassis j is switched on passing from scenario σ to scenario $\sigma + 1$. Finally π_{ij} and $\bar{\pi}$ are parameters representing the power consumption of a single card connecting routers i and j, and the power consumption of a chassis, respectively.

Each traffic demand must be routed along a single path (MPLS routing). This is given by the following flow conservation constraints:

$$\sum_{(i,j) \in A} x_{ij}^{d\sigma} - \sum_{(j,i) \in A} x_{ji}^{d\sigma} = \begin{cases} 1 & \text{if } i = o_d, \\ -1 & \text{if } i = t_d, \\ 0 & \text{otherwise} \end{cases} \quad \forall i \in N, \quad \forall d \in D, \quad \forall \sigma \in S \quad (2)$$

where $x_{ij}^{d\sigma}$ are binary variables that are equal to 1 if the traffic demand d is routed through arc (i, j) in scenario σ. Note that the routing of each demand can be varied along the different time intervals.

There are then the chassis capacity constraints

$$\sum_{(i,j) \in A} \sum_{d \in D} \hat{r}_d^\sigma \rho_d x_{ij}^{d\sigma} + + \sum_{(j,i) \in A} \sum_{d \in D} \hat{r}_d^\sigma \rho_d x_{ji}^{d\sigma} \leq \psi y_j^\sigma, \quad \forall j \in N, \quad \forall \sigma \in S \quad (3)$$

where ψ is the chassis capacity. Demand uncertainty is not taken into account in the node capacity and status constraint. In fact, switching on and off a router is much time and energy consuming than switching on and off a card, and therefore its status changes rarely. Further, many demands are usually routed through on nodes. Thus routers are not likely to switch off due to small variation of demands amount: as long as even one demand is routed through a node, it must be switched on.

The robustness to traffic demand uncertainty is guaranteed by the following robust link capacity constraints:

$$\max_{\substack{\{U_{ij}^{\sigma}\bigcup\{t_{ij}^{\sigma}\}|U_{ij}^{\sigma}\subseteq D,\\ |U_{ij}^{\sigma}|\leq|\Gamma_{ij}^{\sigma}|,\, t_{ij}^{\sigma}\in D\setminus U_{ij}^{\sigma}\}}} \left\{\sum_{d\in U_{ij}^{\sigma}} \bar{r}_d^{\sigma}\rho_d x_{ij}^{d\sigma} + (\Gamma_{ij}^{\sigma}-\lfloor\Gamma_{ij}^{\sigma}\rfloor)\bar{r}_{t_{ij}^{\sigma}}^{\sigma}\rho_d x_{ij}^{t\sigma}\right\} + \\ + \sum_{d\in D} \hat{r}_d^{\sigma}\rho_d x_{ij}^{d\sigma} \leq \mu\gamma w_{ij}^{\sigma}, \quad \forall (i,j)\in A, \quad \forall\sigma\in S \tag{4}$$

The parameter μ is the allowed link capacity utilization. The set $U_{ij}^{\sigma}\bigcup\{t_{ij}^{\sigma}\}$ contains the traffic demands that can be considered uncertain by the link capacity constraint for link (i,j) during scenario σ. Traffic demands belonging to U_{ij}^{σ} can assume the maximal value $\bar{r}_d^{\sigma}$ of their uncertainty set; while $t_{ij}^{\sigma}\notin U_{ij}^{\sigma}$ is an uncertain traffic demand that can assume a maximal value equal to $\left\{\Gamma_{ij}^{\sigma}\right\}\bar{r}_d^{\sigma}$ for capacity constraint on link (i,j) w.r.t. scenario σ.

The power consumed by switching on a chassis is computed through the following constraints:

$$z_j^{\sigma} \geq \delta\bar{\pi}\left(y_j^{\sigma}-y_j^{\sigma-1}\right), \quad \forall j\in N, \quad \forall\sigma\in S \tag{5}$$

where δ is the chassis energy consumption (normalized with respect to hourly chassis consumption) due to a switching-on. We have to keep active the same number of line cards for both the directions of a link:

$$w_{ij}^{\sigma} = w_{ji}^{\sigma}, \quad \forall\sigma\in S, \quad \forall(i,j)\in A: i<j \tag{6}$$

Since, for reasons of reliability, we do not want to switch on a single line card too many times during a single day (too frequent switching can reduce the card life), we added the following constraints to limit to a given ε the allowed maximum number of switching:

$$\sum_{k=1}^{n_{ij}} u_{ijk}^{\sigma} \geq w_{ij}^{\sigma} - w_{ij}^{\sigma-1}, \qquad \forall(i,j)\in A, \forall\sigma\in S \tag{7}$$

$$\sum_{\sigma\in S} u_{ijk}^{\sigma} \leq \varepsilon, \qquad \forall(i,j)\in A, \quad \forall k \tag{8}$$

u_{ijk}^{σ} are auxiliary binary variables which are equal to 1 if cards k-th linking nodes i and j are powered on in scenario σ. For the sake of completeness we also report the domains of the variables:

$$x_{ij}^{d\sigma}\in\{0,1\}, \qquad \forall d\in D, \quad \forall\sigma\in S, \quad \forall(i,j)\in A \tag{9}$$

$$y_j^{\sigma}\in\{0,1\}, \qquad \forall\sigma\in S, \quad \forall j\in N \tag{10}$$

$$z_j^{\sigma}\geq 0, \qquad \forall\sigma\in S, \quad \forall j\in N \tag{11}$$

$$w_{ij}^{\sigma}\in\{0,\ldots,n_{ij}\}, \qquad \forall\sigma\in S, \quad \forall(i,j)\in A \tag{12}$$

$$u_{ijk}^{\sigma}\in\{0,1\}, \qquad \forall\sigma\in S, \quad \forall(i,j)\in A, k\leq n_{ij} \tag{13}$$

3.2 The robust constraints

The robust capacity constraint (4) can be rewritten in a linear form by performing the following calculations: Let

$$\Theta_{ij}^{\sigma} = \max_{\substack{\{U_{ij}^{\sigma}\bigcup\{t_{ij}^{\sigma}\}|U_{ij}^{\sigma}\subseteq D,\\ |U_{ij}^{\sigma}|\leq|\Gamma_{ij}^{\sigma}|,\, t_{ij}^{\sigma}\in D\setminus U_{ij}^{\sigma}\}}} \left\{\sum_{d\in U_{ij}^{\sigma}} \bar{r}_d^{\sigma}\rho_d x_{ij}^{d\sigma} + (\Gamma_{ij}^{\sigma}-\lfloor\Gamma_{ij}^{\sigma}\rfloor)\bar{r}_{t_{ij}^{\sigma}}^{\sigma}\rho_d x_{ij}^{t\sigma}\right\}, \quad \forall\sigma\in S, \quad \forall(i,j)\in A \tag{14}$$

that can be rewritten as

$$\begin{aligned}
\Theta_{ij}^{\sigma} \;=\; & \max_{g} \sum_{d \in D} \bar{r}_d^{\sigma} \rho_d x_{ij}^{d\sigma} g_{ij}^{d\sigma} \\
\text{subject to} \quad & \sum_{d \in D} g_{ij}^{d\sigma} \leq \Gamma_{ij}^{\sigma} \\
& 0 \leq g_{ij}^{d\sigma} \leq \; 1, \qquad \forall d \in D
\end{aligned} \tag{15}$$

where $g_{ij}^{d\sigma} \in [0,1]$ are the real variables that indicate the percentage of maximal variations from the nominal value that is allowed for demand d during scenario σ when link capacity constraint for link (i,j) during scenario σ is considered.

The dual problem of (15) can be expressed as:

$$\begin{aligned}
& \min \sum_{d \in D} l_{ij}^{d\sigma} \; + \; \Gamma_{ij}^{\sigma} \epsilon_{ij}^{\sigma} \\
\text{subject to} \quad & \epsilon_{ij}^{\sigma} + l_{ij}^{d\sigma} \geq \bar{r}_d^{\sigma} \rho_d x_{ij}^{d\sigma} \qquad \forall d \in D \\
& l_{ij}^{d\sigma} \geq 0, \qquad \forall d \in D \\
& \epsilon_{ij}^{\sigma} \geq 0
\end{aligned} \tag{16}$$

We can thus rewrite constraints (4) as:

$$\sum_{d \in D} \hat{r}_d^{\sigma} \rho_d x_{ij}^{d\sigma} + \sum_{d \in D} l_{ij}^{d\sigma} \; + \; \Gamma_{ij}^{\sigma} \epsilon_{ij}^{\sigma} \leq \mu \gamma w_{ij}^{\sigma}, \qquad \forall (i,j) \in A, \quad \forall \sigma \in S \tag{17}$$

$$\epsilon_{ij}^{\sigma} + l_{ij}^{d\sigma} \geq \bar{r}_d^{\sigma} \rho_d x_{ij}^{d\sigma}, \qquad \forall d \in D, \quad \forall (i,j) \in A, \quad \forall \sigma \in S \tag{18}$$

$$l_{ij}^{d\sigma} \geq 0, \qquad \forall d \in D, \quad \forall (i,j) \in A, \quad \forall \sigma \in S \tag{19}$$

$$\epsilon_{ij}^{\sigma} \geq 0, \qquad \forall (i,j) \in A, \quad \forall \sigma \in S \tag{20}$$

where $l_{ij}^{d\sigma}$ and ϵ_{ij}^{σ} are the new variables that let us take into account the supplementary traffic due to the unpredictable variations from the nominal values.

4 The robust heuristic

Since the new robust capacity constraints increase the dimensions of the problem ($|A||S| + |A||S||D|$ new variables and $|A||S||D|$ new constraints), the robust formulation does not allow to efficiently solve at optimality instances with more than 15 nodes and 30 traffic demands. For this reason we have developed a heuristic able to compute solutions with a limited gap from the optimum for instances up to 30 nodes and 400 traffic demands. The procedure is called *Energy-Aware Single Time-period Heuristic* (EA-STH) and manages energy consumption of one time interval at a time and must be repeated for each time interval. The energy consumption of a single time interval is optimized by solving an ILP model: the ILP model is formulated as the one proposed in Section 3.1 but it is applied to only one time interval. When EA-STH is applied to a new time interval, both the impact of chassis switching on and the constraint on the maximum number of card switching on must be taken into account. Thus, suitable parameters are defined, which represent the state of chassis and the number of transitions to on-state for each card in the previous optimized time intervals. All parameters are updated each time a new time interval is optimized, according to the computed solution and they are used in the modified version of constraints (5), (7), (8). To guarantee that constraints on card reliability are not violated, if a card has been already switched on ε times in the previously optimized time intervals, it is forced to keep its current status (powered on) for the next time intervals. Since the final solution can vary according to the starting time period chosen, we repeat the procedure starting the elaboration from all the time intervals and take the best solution

5 Computational results

5.1 The testbed

We have tested and compared the MILP formulation (1)–(13) and EA-STH using the test *9N* network shown in Figure 2 and three networks provided by the SNDLib [22], the *newyork*, *france* and *nobel-eu* networks (see [23] for the figures of the three networks). We assume a power consumption $\bar{\pi}$ equal to $86.4W$ for all the chassis and a power consumption π_{ij} equal to $6.8W$ for all the links. The network nodes are divided into core and edge routers and for each network a subset of core routers was assumed; note that core routers are the only ones that can be put to sleep, since they are neither source nor destination nodes. As for the traffic matrices, the nominal traffic demands ρ_d for SNDLib networks have been obtained by scaling for a fixed parameter ϖ the traffic matrices provided by the SNDLib itself, while as for the test with the 9 nodes network we scaled with ϖ some randomly generated matrices. In both cases ϖ has been dimensioned in order to obtain a utilization lower than 50% (link utilization is usually lower than 50% during peak hours [5]) when nominal demands ρ_d are efficiently routed through the full active networks on single paths. As already mentioned the $\hat{r}_d^\sigma$ parameters (predicted amount of the nominal value ρ_d that has to be satisfied during scenario σ) have been set equal to the average values shown in Figure 1. We have experimented with uncertainty sets for traffic demands of different dimensions and different values of Γ_{ij}^σ (see respectively column $\bar{r}_d^\sigma$ and columns $\Gamma_{ij}^\sigma = 0, 1, 2, 3, 4, 5$ in Tables 1 and 2). Note that in each different test we assume uncertainty sets of the same normalized dimensions (as percentage of $\hat{r}_d^\sigma$) for each demands. We also assume Γ_{ij}^σ parameters equal for each link (i, j) and scenario σ. Finally, we experimented with δ (chassis switching-on normalized consumption) equal to 0.25, ε (switching-on limit) equal to 1, n_{ij} (number of cards in link (i, j)) equal to 2, $\forall (i, j) \in A$) and μ (link max-utilization) equal to 50%.

5.2 The results

The tests have been carried out on Intel i7 processors with 4 core and multi-thread 8x, equipped with 8Gb of RAM. All the computational results are reported in Tables 1 and 2, where $|N|$-$|N_c|$, $|A|$-$|C|$ and $|D|$ represent the number, respectively, of nodes and core nodes, links and line cards and traffic demands. Moreover columns τ^σ and $\bar{r}_d^\sigma$ report respectively, the CPLEX time-limit for solving the single time period formulation, and the dimensions of the uncertainty sets (as percentage of the nominal values $\hat{r}_d^\sigma$). In the following group, $\%E_c$ is the energy consumption level of the optimized network (compared with the consumption of the fully powered on network) and $\%_{Inf}$ is the percentage of random scenarios generated in the uncertainty set that cannot be satisfied by the robust solution found (10000 random scenarios for each instance).

First of all, looking at Table 1, it is interesting to observe that the gap between the solutions obtained by, respectively, solving the robust formulation (1)–(13) and performing EA-STH is very small (equal to 1.4% in the worst case, tests 3–7 with $\Gamma_{ij}^\sigma = 2$) and often equal to 0. The strategy of separately managing

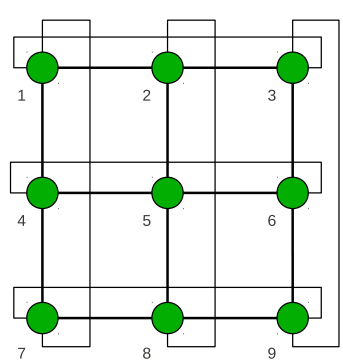

Figure 2: Network with 9 nodes.

the different time interval seems thus promising and valid, since it gives us the possibility to handle bigger networks, providing close to optimal solutions. From Table 1 and 2, we can observe that the energy savings achieved are substantial for all the networks considered and vary from 60% for the *9N* network to 30% for the *france* network with only 7 core nodes (see tests 21–22–23–24). It is very important to note that the cost paid to get solutions robust to traffic variations (by increasing the robustness parameters Γ_{ij}^{σ}) is very limited: the energy gap between the solutions provided by the robust optimization approach and the classical one is generally around 5% for the *9N* network and around 2% for the bigger SNDLib networks. In particular, in the worst case we register an energy consumption increase of 5% (see tests 15–16), while in the best case we assist to an increase of only 0.1% (see test 25). That means that it is possible to obtain more robust solutions by only reorganizing in a more efficient way the demand routing, without the need of reactivating the sleeping devices. The rare cases where a Γ_{ij}^{σ} increase leads to a consumption reduction (see tests 18 and 15) are explained by the gap from the optimum that can vary time to time because of the CPLEX time limit. Note that for the SNDLib networks, it is generally possible to obtain completely robust solutions ($\%_{Inf}$ value close to 0) by setting Γ_{ij}^{σ} parameters equal to 3 or 4. Moreover, as expected, the energy savings are reduced when the uncertainty set dimensions are increased but the reduction is really small. The gap between instances with $\bar{r}_d^{\sigma} = 5\%$ and $\bar{r}_d^{\sigma} = 20\%$ is smaller than 5% for the *9N* network (see the pairs of tests 1–4 and 5–8) and generally around 1% and always smaller than 3% with the SNDLib networks (see the pairs of tests 9–12, 13–16, 17–20, 21–24, 25–28, 29–32).

6 Conclusion

In this paper we considered an energy-aware traffic engineering multi-period problem, and proposed a new approach to add robustness to traffic variations as a key element of the energy-aware optimization. The aim was that of minimizing the energy consumption of an IP network following daily traffic variations, while guaranteeing the satisfaction of all the uncertain traffic demands. The management of uncertain traffic demands is pursued by modifying the capacity constraints in order to account for a limited number of traffic demands different form their nominal values. Our heuristic method called EA-STH, is able to find solutions very close to the optimum with networks of up to 28 nodes. Results show energy savings up to 60% with robust solutions. We are currently working on improving the efficiency of the model and the accuracy of the heuristic, as well as on new inter-periods constraints.

Table 1: Computational results: comparison between the (1)–(13) and the robust heuristic EA-STH with the *9N* test network.

Robust Model						$\Gamma_{ij}^{\sigma}=0$		$\Gamma_{ij}^{\sigma}=1$		$\Gamma_{ij}^{\sigma}=2$	
ID	$\|N\|$-$\|N_c\|$	$\|A\|$-$\|C\|$	$\|D\|$	$\tau^{\sigma}(s)$	$\bar{r}_d^{\sigma}$	$\%E_c$	$\%_{Inf}$	$\%E_c$	$\%_{Inf}$	$\%E_c$	$\%_{Inf}$
1	9-5	36-72	12	/	5%	35.4	98.6	37.9	22.9	37.9	0.0
2	9-5	36-72	12	/	10%	35.4	99.5	38.4	0.0	38.4	0.0
3	9-5	36-72	12	/	15%	35.4	99.9	38.4	24.3	41.7	4.2
4	9-5	36-72	12	/	20%	35.4	99.9	42.0	18.1	41.9	0.0
EA-STH						$\Gamma_{ij}^{\sigma}=0$		$\Gamma_{ij}^{\sigma}=1$		$\Gamma_{ij}^{\sigma}=2$	
ID	$\|N\|$-$\|N_c\|$	$\|A\|$-$\|C\|$	$\|D\|$	$\tau^{\sigma}(s)$	$\bar{r}_d^{\sigma}$	$\%E_c$	$\%_{Inf}$	$\%E_c$	$\%_{Inf}$	$\%E_c$	$\%_{Inf}$
5	9-5	36-72	12	50	5%	35.4	99.6	38.3	18.5	38.4	0.0
6	9-5	36-72	12	50	10%	35.4	99.9	38.4	0.0	38.4	0.0
7	9-5	36-72	12	50	15%	35.4	99.8	38.4	24.0	43.1	3.2
8	9-5	36-72	12	50	20%	35.4	99.9	43.1	18.9	43.3	0.0

Table 2: Computational results with SNDLib networks.

EA-STH							$\Gamma^{\sigma}_{ij}=0$		$\Gamma^{\sigma}_{ij}=1$		$\Gamma^{\sigma}_{ij}=2$		$\Gamma^{\sigma}_{ij}=3$		$\Gamma^{\sigma}_{ij}=4$		$\Gamma^{\sigma}_{ij}=5$	
ID	*Net*	$\lvert N\rvert$-$\lvert N_c\rvert$	$\lvert A\rvert$-$\lvert C\rvert$	$\lvert D\rvert$	$\tau^{\sigma}(m)$	$\bar{r}^{\sigma}_{d}$	$\%E_c$	$\%_{Inf}$	$\%E_c$	$\%_{Inf}$	E_c	$\%_{Inf}$	$\%E_c$	$\%_{Inf}$	$\%E_c$	$\%_{Inf}$	$\%E_c$	$\%_{Inf}$
9	newyork	16-8	98-196	56	10	5%	42.5	97.2	42.6	39.9	42.5	3.4	42.7	0.0	42.6	0.0	42.9	0.0
10	newyork	16-8	98-196	56	10	10%	42.5	99.8	42.7	36.7	42.8	4.4	42.8	0.0	42.9	0.0	43.0	0.0
11	newyork	16-8	98-196	56	10	15%	42.5	100.0	42.7	59.3	43.1	2.7	43.0	0.0	43.1	0.0	43.1	0.1
12	newyork	16-8	98-196	56	10	20%	42.5	100.0	43.0	62.8	43.1	10.1	43.4	0.2	43.6	0.0	43.6	0.0
13	newyork	16-4	98-196	132	10	5%	59.3	100.0	59.8	70.8	60.5	9.0	60.7	1.2	60.8	0.0	61.2	0.1
14	newyork	16-4	98-196	132	10	10%	59.3	100.0	60.8	81.7	61.4	16.2	61.2	0.7	63.6	0.3	63.1	0.0
15	newyork	16-4	98-196	132	10	15%	59.3	100.0	60.8	90.9	61.2	27.9	62.2	3.0	64.7	0.2	63.8	0.2
16	newyork	16-4	98-196	132	30	20%	59.3	100.0	62.0	96.3	64.6	26.1	64.9	4.4	64.9	0.5	64.9	0.1
17	france	25-12	90-180	78	15	5%	56.4	99.0	56.7	51.3	57.3	4.5	57.5	1.3	57.7	0.0	57.8	0.0
18	france	25-12	90-180	78	15	10%	56.4	100.0	57.4	69.3	58.1	12.1	58.0	1.7	57.8	0.0	58.1	0.0
19	france	25-12	90-180	78	15	15%	56.4	100.0	57.9	80.4	58.1	10.1	58.2	1.3	58.8	0.0	58.8	0.0
20	france	25-12	90-180	78	15	20%	56.4	100.0	58.0	88.5	58.7	12.1	59.0	0.6	59.0	0.2	59.2	0.0
21	france	25-7	90-180	153	15	5%	67.6	100.0	67.6	54.1	67.9	20.3	67.9	11.4	68.2	0.0	67.7	0.0
22	france	25-7	90-180	153	15	10%	67.6	100.0	67.7	91.7	68.2	25.2	68.3	2.8	68.3	0.6	68.5	0.0
23	france	25-7	90-180	153	15	15%	67.6	100.0	67.9	93.8	68.1	23.3	68.8	3.5	68.5	0.4	69.6	0.0
24	france	25-7	90-180	153	15	20	67.6	100.0	68.1	97.4	68.6	19.8	69.5	5.1	69.3	0.9	69.4	0.0
25	nobel-eu	28-14	82-164	91	10	5%	59.3	96.9	59.3	49.5	59.3	20.0	59.3	0.4	59.5	0.1	59.4	0.2
26	nobel-eu	28-14	82-164	91	10	10%	59.3	99.9	59.4	48.4	59.7	9.8	59.9	0.5	59.9	0.0	59.8	0.0
27	nobel-eu	28-14	82-164	91	10	15%	59.3	99.9	59.6	55.3	59.7	21.3	59.9	1.1	60.0	0.1	60.1	0.3
28	nobel-eu	28-14	82-164	91	10	20%	59.3	100.0	59.8	73.5	59.9	14.4	60.3	2.0	60.1	0.6	60.3	0.1
29	nobel-eu	28-7	82-164	210	10	5%	65.3	97.8	65.3	72.2	65.4	32.5	65.4	1.4	65.5	1.1	65.5	0.1
30	nobel-eu	28-7	82-164	210	10	10%	65.3	100.0	65.5	75.6	65.5	24.9	65.6	5.1	65.6	0.2	65.7	0.4
31	nobel-eu	28-7	82-164	210	10	15%	65.3	100.0	65.6	75.5	65.7	21.5	65.7	4.7	65.7	0.4	65.8	1.2
32	nobel-eu	28-7	82-164	210	10	20%	65.3	100.0	65.5	85.4	65.6	28.5	65.7	9.9	65.8	0.6	65.9	0.1

References

[1] The Climate Group. SMART 2020: Enabling the low carbon economy in the information age. In *2010 State of Green Business*, June 2008.

[2] J.G. Koomey. Estimating total power consumption by servers in the us and the world, 2007. Technical report, Stanford University, February 2007.

[3] M. Gupta and S. Singh. Greening of the internet. In *Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communications*, pages 19–26, 2003.

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

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

[6] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright. Power awareness in network design and routing. In *Proc. of INFOCOM 08*, pages 457–465, Phoenix, Arizona, 2008.

[7] D. Bertsimas, D.B. Brown, and C. Caramanis. Theory and applications of robust optimization. *Arxiv preprint arXiv:1010.5445*, 2010.

[8] A.P. Bianzino, C. Chaudet, F. Larroca, D. Rossi, and J. Rougier. Energy-aware routing: a reality check. In *GLOBECOM Workshops*, pages 1422–1427. IEEE, 2010.

[9] 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*, 13(2):223–244, 2010.

[10] L. Chiaraviglio, M. Mellia, and F. Neri. Minimizing isp network energy cost: Formulation and solutions. *Networking, IEEE/ACM Transactions on*, PP(99), 2011.

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

[12] E. Amaldi, A. Capone, L.G. Gianoli, and L. Mascetti. Energy management in ip traffic engineering with shortest path routing. In *IEEE WoWMoM 2011, Sustainet Workshop*, 2011.

[13] E. Amaldi, A. Capone, L.G. Gianoli, and L. Mascetti. A milp-based heuristic for energy-aware traffic engineering with shortest path routing. *Network Optimization*, pages 464–477, 2011.

[14] A. Cianfrani, V. Eramo, M. Listanti, and M. Polverini. An ospf enhancement for energy saving in ip networks. In *Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on*, pages 325–330. IEEE, 2011.

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

[16] F. Cuomo, A. Abbagnale, A. Cianfrani, and M. Polverini. Keeping the connectivity and saving the energy in the internet. In *Computer Communications Workshops, INFOCOM*, pages 319–324. IEEE, 2011.

[17] A.P. Bianzino, L. Chiaraviglio, and M. Mellia. Grida: A green distributed algorithm for backbone networks. In *Proc. of GreenCom*, pages 113–119. IEEE, 2011.

[18] E. Gelenbe and T. Mahmoodi. Energy-aware routing in the cognitive packet network. In *ENERGY 2011, The First International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies*, pages 7–12, 2011.

[19] K. Ho and C. Cheung. Green distributed routing protocol for sleep coordination in wired core networks. In *Networked Computing (INC),*, May 2010.

[20] B. Addis, A. Capone, G. Carello, L.G. Gianoli, and B. Sansò. Energy-aware multiperiod traffic engineering with flow-based routing. In *Proc. of ICC 2012, Workshop on Green Communications and Networking, 15th June*. IEEE, 2012.

[21] M. Pióro and D. Medhi. *Routing, Flow and Capacity Design in Communication and Computer Networks*. Morgan Kaufman Publishers, 2004.

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

[23] SNDlib: library of test instances for Survivable fixed telecommunication Network Design. http://sndlib.zib.de/home.action.