Minimization of Network Power Consumption with Redundancy Elimination

Frédéric Giroire[a], Joanna Moulierac[a], Truong Khoa Phan[a,*], Frédéric Roudaut[b]

[a] *Joint project COATI, I3S(CNRS-UNS), INRIA, Sophia-Antipolis, France*
[b] *Orange Labs, Sophia Antipolis, France*

Abstract

Recently, *energy-aware routing* (EAR) has gained an increasing popularity in the networking research community. The idea is that traffic demands are redirected over a subset of the network links, allowing other links to sleep to save energy. In this paper, we propose GreenRE - a new EAR model with the support of *data redundancy elimination (RE)*. This technique, enabled within routers, can virtually increase the capacity of network links. Based on real experiments on Orange Labs platform, we show that performing RE increases the energy consumption for routers. Therefore, it is important to determine which routers should enable RE and which links to put into sleep mode so that the power consumption of the network is minimized. We model the problem as Integer Linear Program and propose greedy heuristic algorithms for large networks. Simulations on several network topologies show that the GreenRE model can gain further 37% of energy savings compared to the *classical EAR* model.

Keywords: Redundancy Elimination, Energy-aware Routing, Algorithm, Green Networking.

*Corresponding author address:
Truong Khoa Phan
INRIA Sophia Antipolis Mediterranee
2004 route des Lucioles, 06902 Sophia Antipolis, France.
Email: truong_khoa.phan@inria.fr
Phone: (+33) (0)4 92 38 71 92
Fax: (+33) (0)4 89 73 24 00

1. Introduction

Recent studies exhibit that traffic load on routers has a small influence on their energy consumption [1, 2]. Instead, the dominating factor is the number of active elements on routers such as ports, line cards, base chassis, etc. The basic idea of energy-aware routing (EAR) is that, during low traffic periods (e.g. at night), traffic demands can be routed over a subset of the network links while preserving connectivity and QoS. In this way, the links excluded by the routing paths can be put into sleep mode (or more precisely, two network interfaces on the two routers will sleep) to save energy.

In general, link capacity is the main constraint of the EAR problem. In this work, we use an assumption that routers can eliminate redundant data traffic and hence, virtually increase capacity of network links. As a result, more traffic flows can be redirected and more links can sleep to save energy. Although routers nowadays cannot remove repeated content from network transfers, there exists WAN Optimization Controller (WOC) - a commercial device used in enterprises or small ISPs to eliminate traffic redundancy [3, 4, 5]. In order to identify the power consumption directly induced by RE, we perform real experiments on the WOC. Because the main idea of routers performing RE is similar to the WOC functionality (see Section 2.2), we believe that when a router eliminates traffic redundancy, it also consumes additional energy like the WOC. In summary, the contributions of this work are the following:

- We do real experiments to exhibit the power consumption of a WOC.
- We define and formulate GreenRE - a new EAR model as Mixed Integer Linear Program (MILP).
- We propose and evaluate a greedy heuristic algorithm that can be used for large-scale networks.
- By simulation, we present energy savings on real network topologies.

The rest of this paper is structured as follows. We summarize related works in Section 2. In Section 3, we model GreenRE as ILP, then propose a greedy heuristic algorithm. Simulation results are presented in Section 4. We present a discussion on EAR in practice in Section 5. Finally, we conclude the work in Section 6.

2. Related Works

2.1. Classical Energy-aware Routing (EAR)

Today's networks are usually built with several redundant links and aggressive over-provision in bandwidth. While these redundancies increase network reliability, they also greatly reduce the network's energy efficiency. Indeed, all network devices are powered on but highly under-utilized most of the time. Since power consumption of a router is independent from its traffic load [1, 2], people proposed to put unused network elements into sleep mode to save energy. This research idea is called *energy-aware routing (EAR)* and is illustrated by the example in Figure 1. Although several works have shown a great opportunity for saving energy using EAR [6, 7], there are a number of issues that arise in practical implementation. We give a discussion on these issues in Section 5.

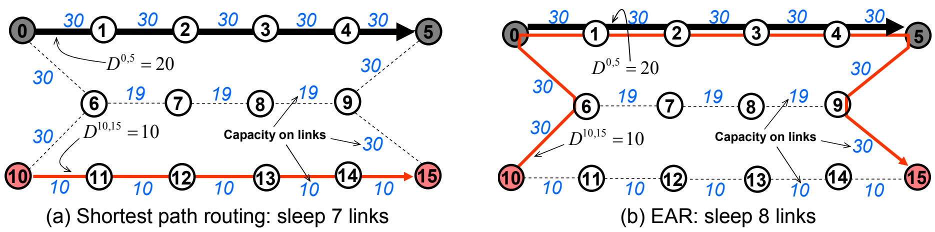

Figure 1: Example of shortest path routing and EAR.

As an example of EAR, we refer to Fig. 1. There are two traffic demands $0 \rightarrow 5$ and $10 \rightarrow 15$ with volumes $D^{0,5} = 20$ Gb and $D^{10,15} = 10$ Gb. The shortest path routing, as shown in Fig. 1a, uses 10 active links whereas the remaining 7 links can be put into sleep mode. However, taking energy consumption into account, in Fig. 1b, EAR solution allows 8 links to sleep, thus energy consumption is further decreased. The problem of minimizing the number of active links under QoS constraints can be precisely formulated using Mixed Integer Linear Programming (MILP). However, this problem is known to be NP-Hard [8], and currently exact solutions can only be found for small networks. Therefore, many heuristic algorithms have been proposed to find admissible solutions for large networks [8, 9].

2.2. Reduction of Traffic Load

Internet traffic exhibits a large amount of redundancy when different users access the same or similar contents. Therefore, several works [10, 11, 12] have explored how to eliminate traffic redundancy on the network. Spring et al. [13] developed the first system to remove redundant bytes from any traffic flows. Following this approach, several commercial vendors have introduced WAN Optimization Controller (WOC) - a device that can remove duplicate content from network transfers [3, 4, 5]. WOCs are installed at individual sites of small ISPs or enterprises to offer end-to-end RE between pairs of sites. As shown in Fig. 2, the patterns of previously sent data are stored in database

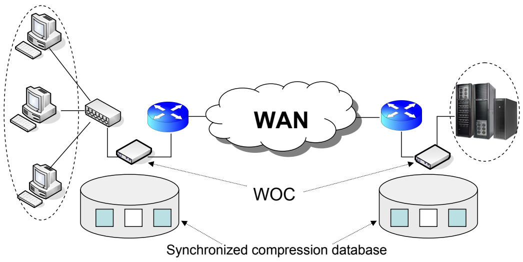

Figure 2: Reduction of end-to-end link load using WOC

of the WOCs at both sending and receiving sides. The technique used to synchronize the databases at peering WOCs can be found in [4]. Whenever the WOC at the sending side notices the same data pattern coming from the sending hosts, it sends a small signature instead of the original data (*encoding process*). The receiving WOC then recovers the original data by looking up the signature in its database (*decoding process*). Because signatures are only a few bytes in size, sending signatures instead of actual data gives significant bandwidth savings.

Recently, the success of WOC deployment has motivated researchers to explore the benefits of deploying RE in routers across the entire Internet [11, 12]. The core techniques used here are similar to those used by the WOC: each router on the network has a local cache to store previously sent data used to encode and decode data packets later on. Obviously, this technique requires heavy computation and large memory for the local cache. However, Anand et al. have shown that on a desktop 2.4 GHz CPU with 1 GB RAM, the prototype can work at 2.2 Gbps for encoding and at 10 Gbps for decoding packets [11]. Moreover, they believe that higher throughput can be attained if the prototype is implemented in hardware. Several real traffic traces have been collected to show that up to 50% of the traffic load can be reduced with RE support [11, 12].

In next section, we propose GreenRE - the first model of *energy-aware* *routing* with RE support. We show that RE, which was initially designed for bandwidth savings, is also potential to reduce network power consumption.

3. Energy-aware Routing with RE

In the GreenRE model, RE is used to virtually increase capacity of the network links. A drawback is that, as shown in [14], when a router performs RE, it consumes more energy than usual. This introduces a tradeoff between enabling RE on routers and putting links into sleep mode. We show that it is a non-trivial task to find which routers should perform RE and which links should sleep to minimize energy consumption for a backbone network.

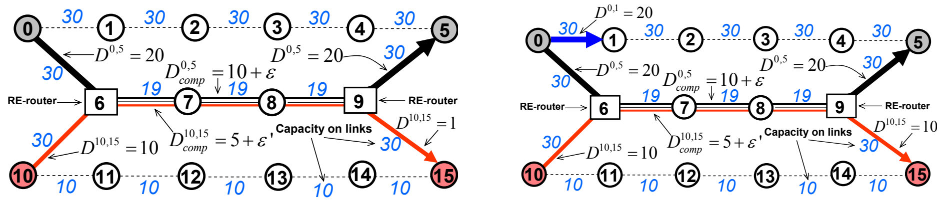

(a) sleep 10 links, enable 2 RE-routers (b) sleep 9 links, enable 2 RE-routers

Figure 3: GreenRE with 50% of traffic redundancy

As an example, we refer to Fig. 3a with two traffic demands $D^{0,5} = 20$ Gb and $D^{10,15} = 10$ Gb. Let a RE-router cost 30 Watts (see Section 4.1) and a link consume 200 Watts [9]. Assume that 50% of the traffic is redundant and RE service is enabled at the router 6 and router 9, thus traffic flows $0 \to 5$ and $10 \to 15$ passing links (6, 7, 8, 9) are reduced to $(10 + \epsilon)$ Gb and $(5 + \epsilon')$ Gb where ϵ, ϵ' denotes the total size of the signatures used for each flow. In reality, each signature is only a few bytes in size [4], therefore ϵ, ϵ' are small and the routing in Fig. 3a is feasible without any congestion. As a result, the GreenRE solution allows to sleep 10 links and enables 2 RE-routers which saves $(10 \times 200 - 2 \times 30) = 1940$ Watts, compared to $8 \times 200 = 1600$ Watts of the EAR solution (Fig. 1b). It is noted that, in some extreme cases, GreenRE even helps to find feasible routing solution meanwhile it is impossible for the classical EAR. For example, if we add a third demand from router 0 to 1 with volume 20 Gb, then Fig. 3b is a feasible solution. However, without RE-routers, no feasible solution is found because there is not enough capacity to route all the three demands.

3.1. Integer Linear Program (ILP) Formulation

The GreenRE model can be formulated as Integer Linear Program. We present a network topology as an undirected graph $G = (V, E)$. The set of nodes V describe routers and the edges $(u, v) \in E$ describe connections between those routers. We note $N(u)$ as a set of neighbor nodes of u in the graph G. For each link $(u, v) \in E$, we use a binary variable x_{uv} to determine if the link is used or not. If link (u, v) is active, two network interfaces at router u and router v are enabled, this consumes PE_{uv} Watts. We define γ^{st} as the percentage of unique (non redundant) traffic. For example, with 40% of redundancy ($\gamma^{st} = 0.6$), instead of sending a traffic demand 10 Gb, we are sending only $(6 + \epsilon)$ Gb after removing redundancy. For simplicity, since ϵ is small, we can ignore it in the formulation and a traffic flow from which redundancy has been removed is called a *compressed flow*. It is noted that, the notion γ^{st} only captures the intra-flow redundancy (and not the inter-flow redundancy as presented in [10]). We note f_{uv}^{st} (resp. g_{uv}^{st}) be the fraction of normal flow (resp. compressed flow) on edge (u, v) corresponding to the demand D^{st} flowing from u to v. We define a binary variable w_u which is equal to 1 if router u performs RE (called RE-router and it consumes additional PN_u Watts). We consider three different scenarios of the problem: (1) all routers on the network can perform RE, we can enable or disable RE service on routers; (2) only a predefined set of routers on the network have RE capability, other routers are normal routers and (3) there is a limited number of RE-routers, the network operators should find where to place them to increase energy efficiency for the network. We formulate the three scenarios of the GreenRE problem as follows:

3.1.1. Scenario 1: All Routers are RE-capable Routers

$$\min \sum_{uv \in E} PE_{uv} x_{uv} + \sum_{u \in V} PN_u w_u \tag{1}$$

$$\text{s.t.} \sum_{v \in N(u)} \left(f_{vu}^{st} + g_{vu}^{st} - f_{uv}^{st} - g_{uv}^{st} \right) = \begin{cases} -1 & \text{if } u = s, \\ 1 & \text{if } u = t, \ \forall u \in V, (s,t) \in D \\ 0 & \text{else} \end{cases} \tag{2}$$

$$\sum_{v \in N(u)} (g_{uv}^{st} - g_{vu}^{st}) \leq w_u \qquad \forall u \in V, (s,t) \in D \tag{3}$$

$$\sum_{v \in N(u)} (g_{vu}^{st} - g_{uv}^{st}) \leq w_u \qquad \forall u \in V, (s,t) \in D \tag{4}$$

$$\sum_{(s,t)\in D} \mathcal{D}^{st}\left(f^{st}_{uv} + f^{st}_{vu} + \gamma^{st}(g^{st}_{uv} + g^{st}_{vu})\right) \leq \mu C_{uv} x_{uv} \quad \forall uv \in E \tag{5}$$

$$f^{st}_{uv}, \quad g^{st}_{uv} \in \ [0,1], \qquad w_u, \quad x_{uv} \in \{0,1\} \tag{6}$$

The objective function (1) is to minimize the power consumption of the network represented by the number of active links and RE-routers. Constraints (2) establish flow conservation constraints. Constraints (3)-(4) are used to determine whether RE service is enabled at router u or not. If it is not ($w_u = 0$), the router u only forwards flows without compression or decompression, then the amount of compressed flows incoming and outgoing the router u is unchanged. It is noted that if a flow is compressed, it needs to be decompressed somewhere on the way to its destination. This requirement is implicitly embedded in the constraints (4). For instance, assume that a destination node t is not a RE-router ($w_t = 0$). When a compressed flow g^{st} reaches its destination, because t is the last node on its path, the flow can not be decompressed. Consider the constraints (4), we have $u = t$, then $\sum_{v\in N(u)} g^{st}_{vt} > 0$ (the compressed flow enters node t) and $\sum_{v\in N(u)} g^{st}_{tv} = 0$ (t is the destination node). Therefore, the constraint (4) is violated and the flow should be decompressed before or at least at the destination node ($w_t = 1$). We consider an undirected link capacity model [15] in which the capacity of a link is shared between the traffic in both directions. We use constraints (5), where μ denotes the link utilization in percentage, to limit the available capacity of a link.

3.1.2. Scenario 2: a Predefined Set of Routers are RE-capable Routers

We define the following constraints:

$$w_u = 0 \qquad\qquad \forall u \notin V', \quad V' \subset V, \tag{7}$$

where V' is a predefined subset of routers that have RE-capability, we force all other routers to be normal routers ($w_u = 0$). By adding (7) to the first scenario (1)–(6), we have the second scenario of the GreenRE problem.

3.1.3. Scenario 3: a Limited Numbers of RE-routers

We add the following constraints to the first scenario (1)–(6):

$$\sum_{u\in V} w_u \leq M, \tag{8}$$

where M is a parameter denoting the maximum number of RE-routers that can be placed on the network. By using constraints (8), we allow any router to perform RE. However, the total number of RE-routers on the network should be less than M.

3.2. Heuristic Algorithm

Energy-aware routing problem is known to be NP-Hard [8, 16]. We have presented a greedy heuristic called *H-GreenRE* for large network topologies in our previous work [14]. Since the power consumption of a link is much more compared to an enabled RE-router, the heuristic should give priority to minimize the number of active links. The basic idea of *H-GreenRE* is that we first assume all routers are RE-enabled routers. Then, the traffic flows can be compressed everywhere on the network. Based on this, we find feasible routing solutions (without overloaded links) and then turn off the link with minimum traffic load. This procedure is repeated until there is no link that can be turned off. After that, we fix the routing and based on some rules defined in [14], we disable unnecessary RE service on routers to save energy. The main problem of *H-GreenRE* is that, to find a feasible routing, we pick up demand one by one and try to route it using shortest path routing. Thus, in case the feasible routing is not a shortest path, *H-GreenRE* can not find it. In this paper, we introduce a new heuristic algorithm based on the ILP formulation (let's call it H_{ILP}-*GreenRE*). Using the ILP formulation, H_{ILP}-*GreenRE* tries all the possibilities to find a feasible routing if any, thus it is more efficient than the *H-GreenRE* (see Section 4.2.1). In the GreenRE problem, to find optimal solution we must consider power consumption of both active links and RE-routers at the same time. For the heuristic H_{ILP}-*GreenRE*, we divide the original problem into two sub-problems: (1) minimize the number of active links and (2) minimize the number of RE-routers. Similar to the *H-GreenRE*, the H_{ILP}-*GreenRE* has two steps: the first step is to find as few active links as possible (sub-problem 1), and then we minimize the number of RE-routers in the second step (sub-problem 2).

Step 1 (sub-problem 1) is a constraint satisfaction problem returning a feasible routing solution. We use the same framework for the three scenarios of the GreenRE problem. For details, to find feasible solutions ($P_{current}$ - line 2 and P_{new} - line 8), we set the objective function to *min* 0 and use the constraints (2)–(6) for scenario 1. Similarly, scenario 2 (resp. scenario 3) uses the constraints (2)–(7) (resp. constraints (2)–(6), (8)) and the objective is *min* 0. In each round of the algorithm, we try to remove a link with low

Algorithm 1: *Inputs:* A graph $G = (V, E)$ with link capacity C_e, a set of traffic demands and non-redundant rates.

1 *Step 1 - Minimize number of active links by removing low loaded links:*
2 Find a feasible routing solution using the ILP called *P_current*;
3 Let S be an ordered list initialized with the links of G sorted by increasing traffic load in *P_current*;
4 Let $R := \emptyset$ be the set of links that cannot be removed;
5 **repeat**
6     $e := S.lowest\_loaded\_link()$ such that $e \notin R$;
7     $S := S \backslash \{e\}$;
8   **if** *a feasible routing* P_new *on* $E \backslash \{e\}$ *is found* **then**
9   **if** *P_new has less active links than P_current* **then**
10   *P_current* := *P_new*;
11             $S :=$ list of links sorted by increasing traffic load in *P_new*;
12             $E := E \backslash \{e\}$;
13   **end**
14   **else**
15         $R := R \cup \{e\}$;
16   **end**
17 **until** $(S = \emptyset)$ *or* $(R = S)$;
18 Return the final feasible routing solution (if any);
19 *Step 2* - Find feasible solution minimizing the number of RE-routers on the set of active links E found in *Step 1*.

load (line 6 - 7) and then find a new feasible routing (P_{new} - line 8) using less active links. The idea behind this algorithm is that we try to put into sleep mode the low loaded links and to accommodate their traffic on other links in order to reduce the total number of active links. Observe that unused links (i.e. links that do not carry traffic) are not considered in the set S since the removal of such a link will result in a routing *P_new* equal to *P_current*. To further reduce the computation time, we can consider additional heuristic. For instance, while removing a low loaded link (line 6 - 7), we can also set the variable x_{uv} associated to a heavily loaded link to 1 so that it can speed up the resolution for finding P_{new} (line 8). Indeed, such high loaded link will certainly be part of the final solution. Since we relax the objective function and the goal is just to justify whether a set of constraints is feasible or not, it is quite fast to find $P_{current}$ and P_{new}. In our simulations, the execution time of Algorithm 1 (including the two steps) is less than one hour for the tested network topologies (see Section 4).

After *Step 1*, if a feasible routing is found, and so a set of active links, we proceed to *Step 2* (sub-problem 2) to minimize the number of enabled RE-routers. More precisely, we use again the ILP formulation (of the scenario we want to solve) in which the objective function is set to $\min \sum_{u \in V} w_u$. Furthermore, we set all binary variables associated to active links to 1 and the others to 0 (this speed-up the resolution of the ILP).

4. Experiment and Simulation Results

4.1. Energy Consumption with WOC

Several results of bandwidth savings using WOC can be found in [4]. We have also performed experiments on the network platform of the project Network Boost at Orange Labs (the full figure of the test-bed can be found in [17]). We installed two WOCs, each at the access link of the two sites (let's call them site A and site B). These two sites are connected via a backbone composed of 4 routers. We setup FTP connections for uploading files from site A to site B. As shown in Fig. 4a, power consumption of the WOC is

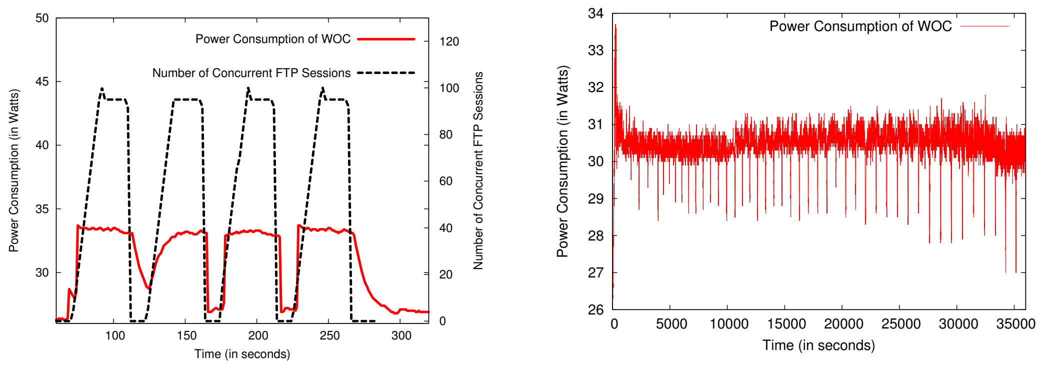

(a) With concurrent FTP sessions

(b) With increasing size of uploaded files

Figure 4: Power consumption of the WOC

increased (from 26 Watts to 34 Watts) with the number of concurrent FTP sessions. For the next experiment, we keep only one FTP session and let the WOC perform RE for 10 hours in which the sizes of uploaded files are increased. The results show that the WOC consumes around 30 Watts on average (Fig. 4b). Therefore, for sake of simplicity, we use an average value of power consumption (30 Watts) to represent additional cost for the router to perform RE.

4.2. Simulation Results with GreenRE

We solve the GreenRE model with IBM CPLEX 12.4 solver [18]. All computations were carried out on a 2.7 Ghz Intel Core i7 with 8 GB RAM. We studied ten classical real network topologies extracted from SNDLib [19]. Their sizes span from 15 to 54 nodes and from 22 to 89 edges, as summarized in Table 1. According to the results of the works mentioned in Section 2.2, we use redundancy rates equal to 50% (high redundancy, $\gamma = 50\%$), 25% (medium redundancy, $\gamma = 75\%$) and 10% (low redundancy, $\gamma = 90\%$). For worst-case scenario and for comparison with previous work [8, 14], all links are set up with the same capacity C and the demands are all-to-all (one router has to send traffic to all remaining routers on the network) with the same traffic volume D for each demand.

4.2.1. Comparison with the Heuristic H-GreenRE in [14]

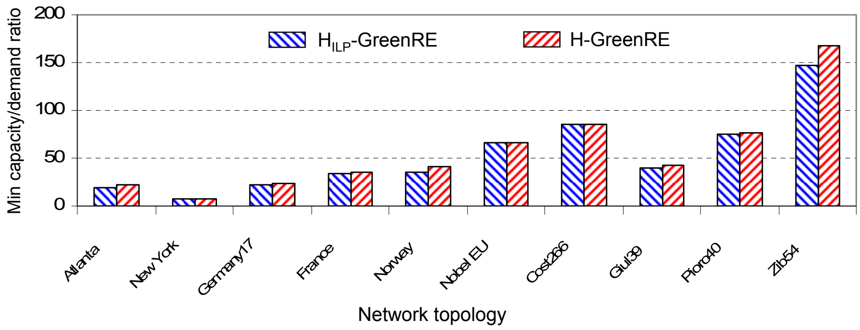

Figure 5: Comparison of λ_{minRE} between H_{ILP}-GreenRE and H-GreenRE

We propose in this paper a new heuristic based on the ILP formulation called *H_{ILP}-GreenRE*. To compare with the heuristic *H-GreenRE* proposed in the previous work [14], two simulation scenarios (with $\gamma = 50\%$) have been done for the ten network topologies (we sort the networks in increasing order of the number of nodes).

First, we find the minimum values of capacity/demand ratio λ_{minRE} that allow for each heuristic algorithm to find a feasible routing solution *with the support of RE-routers*. Note that, λ represents the level of traffic load on

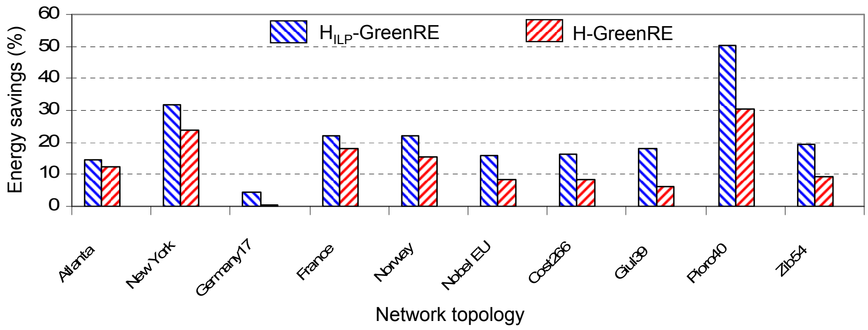

Figure 6: Comparison of energy savings between H_{ILP}-GreenRE and H-GreenRE

the network. Small value of λ means that the traffic load on the network is high (e.g. traffic at peak hours), thus it is hard to find feasible solution because of the lack of capacity (refer to the example in Fig. 3b). Therefore, the heuristic algorithm that can find feasible routing with smaller value of λ_{minRE} is the better one. To compute λ_{minRE}, we first fix the demand value, e.g. $D = 1$. Then, starting with a large capacity value, e.g. $C = 1000$, we decrease the value of C and test the heuristic until we get the minimum value of C that is still possible to find a feasible solution. Let's call this value is C_{min}, then we have $\lambda_{minRE} = C_{min}$. Fig. 5 shows that H_{ILP}-GreenRE can find feasible solutions with smaller values of λ_{minRE} than H-GreenRE. For example, for the Atlanta network, H_{ILP}-GreenRE finds a solution with $\lambda_{minRE} = 19$ while H-GreenRE is with $\lambda_{minRE} = 22$: that is, for example, for a link capacity of 10 Gbit/sec, the first heuristic succeeds in routing an all-to-all demand of 10/19 =0.53 Gbit/sec for each demand and the second heuristic, a demand of only 10/22 = 0.45 Gbit/sec. In summary, H_{ILP}-GreenRE finds feasible solutions close to the lower bounds of λ_{min} found in [8]. The best improvement is on Zib54 network: $\lambda_{minRE} = 147$ (for H_{ILP}-GreenRE) in comparison with $\lambda_{minRE} = 168$ (for H-GreenRE).

We show next the energy savings for the ten networks. We use the value of λ_{minRE} that allows for H-GreenRE to find feasible routing solution for each network (the second column in Fig. 5). If a network has dense links, there are more chances to redirect traffic and put links into sleep mode, thus more energy can be saved. As shown in Fig. 6, H_{ILP}-GreenRE again outperforms H-GreenRE for all the networks. Energy efficiency can be increased from 2% (Atlanta network) to 19.8% (Pioro40 network).

4.2.2. Energy Savings for Atlanta Network

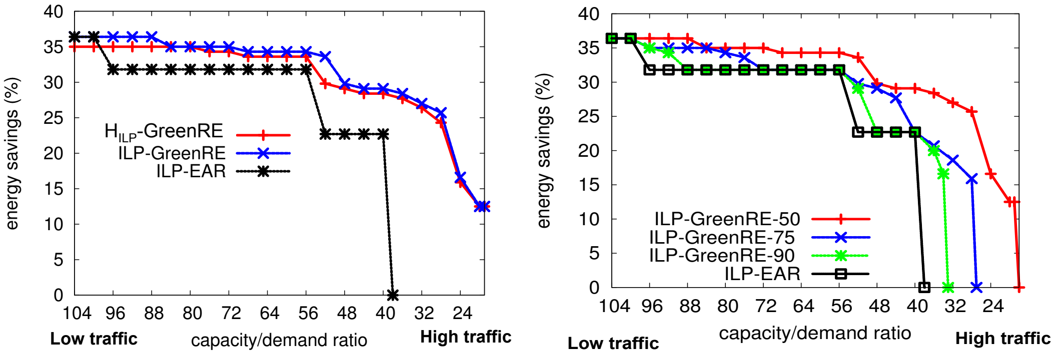

(a) Heuristic vs. ILP-GreenRE vs. ILP-EAR

(b) with different RE rates

Figure 7: Simulation results for Atlanta network

In this subsection, we present simulation results for Atlanta network. In Fig. 7a, with the same redundancy rate ($\gamma = 50\%$), we vary capacity/demand ratios and compare between the *ILP-EAR without* RE-routers (given by [8]), the ILP with RE-routers (the formulation (1) - (6): *ILP-GreenRE*) and the heuristic with RE-routers (H_{ILP}-*GreenRE*). Even for small network like Atlanta, CPLEX also takes some hours to find an optimal solution when the capacity/demand ratios are high (e.g. $\lambda \geq 48$). It is noted that when $\lambda < 48$, it is possible to find an optimal solution within one hour. We limit the solving time to one hour for all instances of Atlanta network corresponding to different capacity/demand ratios. In average, the optimality gap is within 10% for all the best solutions. The heuristic is quite fast, it takes less than 10 seconds to find a solution. The $x-$axis in Fig. 7a represents the capacity/demand ratio λ and the $y-$axis is energy savings in percentage. As shown in Fig. 7a, without RE-router (ILP-EAR), there is no feasible routing solution and hence, no energy is saved if $\lambda < 38$. When λ increases, links have more bandwidth to aggregate traffic, the solutions with and without RE-router converge to the same amount of energy savings. In general, the heuristic with RE-routers works well and approximates to the results of ILP-GreenRE (the max gap is 3.8%).

In Fig. 7b, we evaluate energy savings for Atlanta network with different level of redundancy. It is clear that when traffic redundancy is high, e.g. $\gamma = 50\%$, more traffic flows are aggregated and thus, more links can be turned off to save energy. Similarly, when $\gamma = 75\%$ and $\gamma = 90\%$ (corresponding to 25% and 10% of traffic redundancy), less energy can be saved. These remarks can be seen in Fig. 7b where the gaps between *ILP-GreenRE* and *ILP-EAR* are reducing when γ is increasing. It is noted that *ILP-GreenRE* should be at least as good as *ILP-EAR*. It is because the objective of *ILP-GreenRE* is to minimize energy consumption for the network. In case the redundancy elimination does not help to turn off more links, *ILP-GreenRE* does not enable RE service on router (even it helps to reduce the traffic load). Therefore, in the worst case scenario (redundancy rate is zero or $\gamma = 100\%$), *ILP-GreenRE* has no RE-enabled router and the routing solution is the same as in *ILP-EAR*.

4.2.3. Energy Savings for the Ten Classical Networks

Network	$\mid V \mid$	$\mid E \mid$	λ_{min}	Traffic volume (capacity/demand ratio λ)					
				with RE-router			without RE-router		
				λ_{min}	$2\lambda_{min}$	$3\lambda_{min}$	λ_{min}	$2\lambda_{min}$	$3\lambda_{min}$
Atlanta	15	22	38	27.7%	34.3%	36.4%	0%	32%	36%
New York	16	49	15	52.2%	62.9%	65.8%	2%	59%	63%
Germany17	17	26	44	30.6%	36.7%	37.3%	0%	35%	39%
France	25	45	67	39.2%	43.4%	46%	0%	42%	44%
Norway	27	51	75	37.7%	45.6%	47.8%	12%	43%	47%
Nobel EU	28	41	131	29.2%	33.1%	34.2%	12%	32%	34%
Cost266	37	57	175	30.6%	35%	36.3%	3.5%	32%	35%
Giul39	39	86	85	42.5%	50.5%	53.3%	0%	45%	50%
Pioro40	40	89	153	50.5%	53.7%	55.2%	0%	53%	54%
Zib54	54	80	294	27.5%	30.8%	32.8%	0%	30%	33%

Table 1: Gain of energy consumption (in %)

We present in Table 1 energy gain for ten classical network topologies using H_{ILP}-GreenRE and H-EAR - the heuristic *without* RE-routers found in [8]. Different from λ_{minRE} in Section 4.2.1, we use λ_{min} be the smallest value of capacity/demand ratio that allows to find a feasible route for all the demands *without RE-router* (found in [8]). In the simulations, a range of $\lambda = \{\lambda_{min}, 2\lambda_{min}, 3\lambda_{min}\}$ is used to represent high (e.g. traffic at peak hours), medium and low traffic load (e.g. traffic at night) on the networks. As shown in Table 1, with RE-routers, it starts to save a large amount of energy (in average 37%) even with $\lambda = \lambda_{min}$. Recall that routing with RE-routers is possible even with $\lambda < \lambda_{min}$ meanwhile no feasible solution is found without RE-router. When λ is large enough, it is not necessary to have RE-routers on the network, therefore both the solutions (with and without RE-router) converge to almost the same value of gains in energy savings.

4.2.4. Energy Savings for Scenario 2 and Scenario 3 of the GreenRE Problem

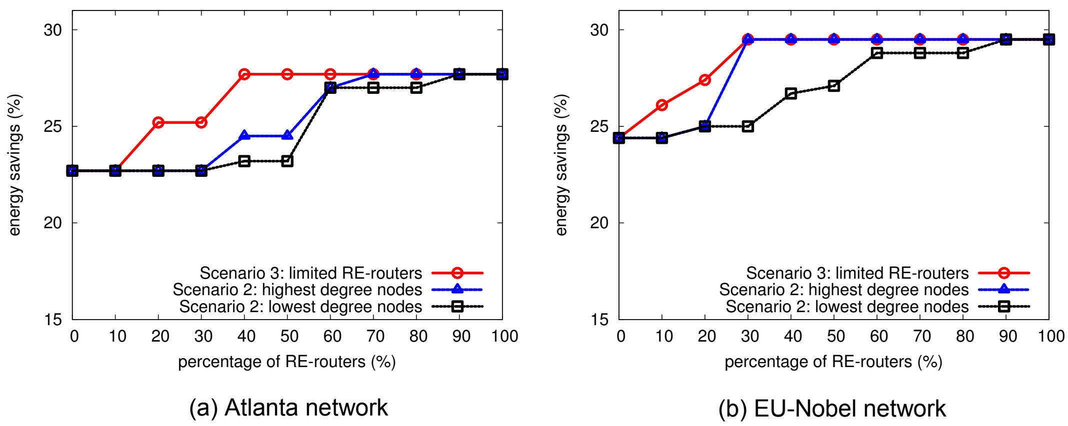

Figure 8: Energy savings with limited RE-routers vs. a subset of capable RE-routers

In this section, we evaluate energy savings of scenario 2 (a predefined subset of RE-capable routers) and scenario 3 (a limited numbers of RE-capable routers). We set link capacity and demand corresponding to λ_{min} in Section 4.2.3. The x-axis of Fig. 8 is the percentage of RE-capable routers on the network. For instance, with scenario 3, we find the routing solution that minimizes energy consumption while there are at most $(x \times |V|)$ RE-routers on the network. For scenario 2, we place $(x \times |V|)$ RE-capable routers on (1) highest degree nodes or (2) lowest degree nodes in graph G. As shown in Fig. 8, the scenario 3 always outperforms the scenario 2 since it can find best positions to place RE-capable routers. For instance, in Atlanta network with a maximum of 6 RE-routers, the max gap is 4.5% and there are 4 RE-routers at the highest degree nodes and the two others are at the medium and the lowest degree nodes. Another important observation we found in the scenario 2 is that, placing RE-routers on high degree nodes gives better results in energy savings. It is because placing RE-capable routers on high degree nodes helps to reduce traffic load and gives more chances to redirect traffic on a few links, allowing other links on these nodes to sleep.

5. Discussion

While the results in this work show a great improvement on energy efficiency for backbone networks, there are still a number of issues of EAR in practice that we discuss below. Nowadays, network operators do not like to turn-off links and change routing. As a matter of fact,

- the convergence between two routing configurations is usually long and can take up to several seconds for some distributed protocols (e.g. for BGP), leading to perturbations for the network traffic.
- even signaling takes non-negligible time: a simple control message may take tens of milliseconds to traverse the network and reach a given router.

However, in our study, we are in a simplified configuration as we consider:

- Preplanned changes: our approach in this paper is completely off-line as we are studying a single traffic matrix. We are not in a context of rapid changes for the routing. Considering a given network and an estimation of the traffic, we can precompute a set of configurations depending on different time instants for a whole day. The number of configurations depends on several parameters that need to be discussed with the network operator and optimized, e.g., link over-dimensioning, QoS, length of paths, number of desired routing changes per day, etc. For instance, as shown in [20], only few configurations (two or three per day) are enough to significantly save energy. Similarly, the authors in [21] suggest to use only two network configurations for a daily time. During peak hours, all the links in the network are active while EAR (with sleep links) is considered only in off-peak time to save energy.
- Intra-domain network: our work focus on intra-domain network, which is a simpler case as the routing changes do not need to be propagated to other domains by inter-domain routing protocols such as BGP. Moreover, the network operator can manage its own domain without explicit cooperation with other domains and update metrics for the routing protocol, turn on/off links, or set up a centralized entity in order to perform a better optimization with a complete view of the network and the traffic.

- Centralized control (SDN): within a network operator domain, an SDN solution, like Openflow, can be implemented. OpenFlow is a promising method to implement EAR in a network [22, 23]. Without setting entries manually, a centralized controller can collect traffic matrices, perform routing calculation in advance and then trigger an installation of new routing rules on routers at the precomputed time based on the planification. Openflow will make energy-aware solutions possible by allowing to switch on and off equipments based on traffic measures. Such a system will experience a much lower routing convergence time than existing distributed protocols [24].

Even if our work focus on a simplified version of the whole problem, we discuss below the problems that may arise when passing from a routing configuration to another, or if unexpected events appear.

- *Transition between two configurations.* We must ensure that the network behaves correctly between two consecutive transitions. In [24], the authors describe a mechanism using Openflow for consistent network updates in which the packets are processed either with the old or the new configuration but never with a mixture of the two. One of the ideas presented in the paper is to pre-install the new routing rules while keeping the old routing in place. When all flows matching a given old rule finish, the new routing configuration takes effect. In [25], the authors use OSPF routing protocol and propose incremental updates of links (turning on and off links is also considered) in order to ensure a loop-free transition between two configurations. In [26], we proposed methods to limit OSPF changes between two consecutive routing configurations to reduce network oscillation during transient time. Moreover, as shown in [27], an appropriate scheduling and prioritization in rule update can help to reduce the side effects in transitioning between network configurations.

- *Unexpected events.* In case of traffic burstiness, the network should react quickly in order to avoid packet losses. Similarly to what is actually done by telecommunication operators, our model integrates over provisioning on links (parameter μ in the linear program). This allows to tolerate some traffic increase without changing the network configuration. If the capacity on links is not sufficient enough, then the SDN controller can send signals to wake up all the links and apply the routing configuration for peak hours.

In this case, the delay will include the time for the controller to contact all the routers, and the routers to change their routing configurations. This will be the worst-case scenario implying the longer delay.

- *QoS degradation and routing loops.* Another issue of EAR is QoS degradation during transient times such as routing loops. Moreover, during these periods, packets may arrive out of order which causes problems for TCP, degrading the perceived QoS for end-users [28]. These effects can be limited if two consecutive routing configurations contain a minimal number of changes. Moreover, we believe that these periods will be short (in the order of seconds) and will be acceptable in view of the gain in energy savings.

- *Delay for (de)/activating links.* The time required for sleeping/activating a link on router is substantial. Indeed, commercial routers available today do not integrate this sleep mode functionality yet. However, there are on-going works on this issue [29]. For instance, the works in [30, 31] show that it is possible to put a link into sleep mode and wake it up in short time (e.g. less than $5\mu s$ for 10GBASE-T links). Thus, we can expect that these advances may come in the future, especially if they offer big energy savings, and will be negligeable. However, this delay is not only for the material to be turned-on but also for the controller to contact all the routers, and the routers to change their routing table.

To sum up, turning-off links in EAR will cause unavoidable overhead, especially in transient time between different routing configurations. However, if the planification of these changes is done with a reasonable number of configurations per day, with few changes between two configurations in order to limit the routing loops and the network oscillations, and with an ordered list of routing changes to be sent to the routers by the SDN controller, we can expect that these drawbacks would be acceptable in view of the energy savings.

6. Conclusion

To the best of our knowledge, GreenRE is the first work considering redundancy elimination as a complementary help for energy-aware routing problem. We formulate the problem as Mixed Integer Linear Program and propose greedy heuristic algorithms. The simulations on several network topologies show a significant gain in energy savings with GreenRE. For future work, we will consider a more realistic model in which data redundancy rates and traffic demand volumes variate based on real life traffic traces. Moreover, we plan to study the inter-flow redundancy as it could further reduce network traffic.

Acknowledgment

This work has been done in collaboration with Network Boost Project (Orange Labs) and partly funded by the ANR DIMAGREEN. The authors would like to thank Didier Leroy (Orange Labs), Yaning Liu (JCP-Consult R&D) and David Coudert (Project COATI) for their advices and support.

References

[1] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, S. Wright, "Power Awareness in Network Design and Routing", in: IEEE International Conference on Computer Communications (INFOCOM), 2008.

[2] P. Mahadevan, P. Sharma, S. Banerjee, "A Power Benchmarking Framework for Network Devices", in: International Conferences on Networking (IFIP NETWORKING), 2009, pp. 795–808.

[3] BlueCoat: WAN Optimization, `http://www.bluecoat.com/`.

[4] T. J. Grevers, J. Christner, "Application Acceleration and WAN Optimization Fundamentals", in: Cisco Press, 2007.

[5] `http://www.riverbed.com/us/solutions/wan_optimization/`.

[6] R. Bolla, R. Bruschi, F. Davoli, F. Cucchietti, "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures", IEEE Communication Surveys and Tutorials 13 (2011) 223 – 244.

[7] A. P. Bianzino, C. Chaudet, D. Rossi, J. Rougier, "A Survey of Green Networking Research", IEEE Communication Surveys and Tutorials 14 (2012) 3 – 20.

[8] F. Giroire, D. Mazauric, J. Moulierac, B. Onfroy, "Minimizing Routing Energy Consumption: from Theoretical to Practical Results", in: IEEE/ACM Green Computing and Communications (GreenCom), 2010, pp. 252–259.

[9] L. Chiaraviglio, M. Mellia, F. Neri, "Minimizing ISP Network Energy Cost: Formulation and Solutions", IEEE/ACM Transaction in Networking 20 (2011) 463 – 476.

[10] A. Anand, A. Gupta, A. Akella, S. Seshan, S. Shenker, “Packet Caches on Routers: the Implications of Universal Redundant Traffic Elimination”, in: ACM Special Interest Group on Data Communication (SIGCOMM), 2008, pp. 219–230.

[11] A. Anand, V. Sekar, A. Akella, “SmartRE: an Architecture for Coordinated Network-wide Redundancy Elimination”, in: ACM Special Interest Group on Data Communication (SIGCOMM), 2009, pp. 87–98.

[12] Y. Song, K. Guo, L. Gao, “Redundancy-aware Routing with Limited Resources”, in: International Conference on Computer Communications and Networks (ICCCN), 2010, pp. 1–6.

[13] N. T. Spring, D. Wetherall, “A Protocol-Independent Technique for Eliminating Redundant Network Traffic”, in: ACM Special Interest Group on Data Communication (SIGCOMM), 2000.

[14] F. Giroire, J. Moulierac, T. K. Phan, F. Roudaut, “Minimization of Network Power Consumption with Redundancy Elimination”, in: International Conferences on Networking (IFIP NETWORKING), 2012, pp. 247–258.

[15] C. Raack, A. M. C. A. Koster, S. Orlowski, R. Wessäly, “On Cut-based Inequalities for Capacitated Network Design Polyhedra”, Networks 57 (2011) 141 – 156.

[16] A. Koster, T. K. Phan, M. Tieves, “Extended Cutset Inequalities for the Network Power Consumption Problem”, in: International Network Optimization Conference (INOC), Electronic Notes in Discrete Mathematics, 2013, pp. 69–76.

[17] www-sop.inria.fr/mascotte/rapports_stages/KhoaPhan_internship-2011.pdf.

[18] IBM ILOG, CPLEX Optimization Studio 12.4.

[19] S. Orlowski, R. Wessäly, M. Pióro, A. Tomaszewski, SNDlib 1.0 - survivable network design library, Networks 55 (3) (2010) 276–286.
URL http://sndlib.zib.de

[20] L. Chiaraviglio, A. Cianfrani, E. L. Rouzic, M. Polverini, “Sleep Modes Effectiveness in Backbone Networks with Limited Configurations”, Computer Networks 57 (2013) 2931–2948.

[21] F. Francois, N. Wang, K. Moessner, S. Georgoulas, K. Xu, “On IGP Link Weight Optimization for joint Energy Efficiency and Load Balancing Improvement”, Computer Communications.

[22] X. Wang, Y. Yao, X. Wang, K. Lu, Q. Cao, "CARPO: Correlation-Aware Power Optimization in Data Center Networks", in: Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), 2012.

[23] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, N. McKeown, "ElasticTree: Saving Energy in Data Center Networks", in: USENIX conference on Networked systems design and implementation (NSDI), 2010.

[24] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, D. Wamker, Abstractions for network update, in: ACM Special Interest Group on Data Communication (SIGCOMM), 2012.

[25] P. Francois, M. Shand, O. Bonaventure, Disruption free topology reconfiguration in ospf networks, IEEE INFOCOM 2007.

[26] J. Moulierac, T. K. Phan, "Optimizing IGP Link Weights for Energy-efficiency in a Changing World", `http://hal.inria.fr/hal-00988882`, iNRIA Research Report (2014).

[27] P. Peresini, M. Kuzniar, M. Canini, D. Kostic, "ESPRES: Easy Scheduling and Prioritization for SDN", in: Open Networking Summit, 2014.

[28] K. C. Leung, V. O. K. Li, D. Yang, "An Overview of Packet Reordering in Transmission Control Protocol (TCP): Problems, Solutions, and Challenges", IEEE Transactions on Parallel and Distributed Systems 18 (2007) 522 – 535.

[29] Econet project, `http://www.econet-project.eu`.

[30] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, J. A. Maestro, "IEEE 802.3az: The Road to Energy Efficient Ethernet", IEEE Communications Magazine 48 (2010) 50–56.

[31] R. Bolla, R. Bruschi, A. Carrega, F. Davoli, P. Lago, "A Closed-Form Model for the IEEE 802.3az Network and Power Performance", IEEE Journal on Selected Areas in Communications 32 (2014) 16 – 27.