



HAL
open science

Minimizing Routing Energy Consumption

Frédéric Giroire, Joanna Moulierac, Brice Onfroy

► **To cite this version:**

Frédéric Giroire, Joanna Moulierac, Brice Onfroy. Minimizing Routing Energy Consumption. [Research Report] RR-7234, 2010. inria-00464318v1

HAL Id: inria-00464318

<https://inria.hal.science/inria-00464318v1>

Submitted on 16 Mar 2010 (v1), last revised 10 Feb 2011 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Minimizing Routing Energy Consumption

Frédéric Giroire — Joanna Moulhierac — Brice Onfroy

N° 7234

Mars 2010

Thème COM



*R*apport
de recherche



Minimizing Routing Energy Consumption

Frédéric Giroire * , Joanna Moulrierac * , Brice Onfroy *

Thème COM — Systèmes communicants
Projet Mascotte

Rapport de recherche n° 7234 — Mars 2010 — 16 pages

Abstract: Several studies exhibit that the traffic load of the routers only has a small influence on their energy consumption. Hence, the power consumption in networks is strongly related to the number of active network elements, such as interfaces, line cards, base chassis,... The goal thus is to find a routing that minimizes the (weighted) number of active network elements used when routing. In this paper, we consider a simplified architecture where a connection between two routers is represented as a link joining two network equipments. When a connection is not used, both network equipments can be turned off. Therefore, in order to reduce power consumption, the goal is to find the routing that minimizes the number of used links while satisfying all the demands. We first present a study on specific topologies, such as trees and complete graphs, that provide bounds and results useful for real topologies. We model the problem as a linear program and propose a heuristic to solve large instances. We exhibit the gain in terms of number of network equipments (leading to a global reduction of the power consumption) for a set of network topologies: we see that for almost all topologies more than one third of the network equipments can be spared for usual ranges of operation. Finally, we discuss the impact of energy efficient routing on the stretch factor and on fault tolerance.

Key-words: power consumption, energy-efficient routing, graphs, linear programming.

Mail: firstname.name@inria.fr

This work has been partly funded by the European project FET AEOLUS, the ANR JCJC DIMA-GREEN.

* MASCOTTE, INRIA, I3S (CNRS - University Nice Sophia-Antipolis), France

Routage efficace en énergie

Résumé : L'économie d'énergie dans les réseaux peut être accomplie en utilisant des techniques efficaces de routage ou de conception de réseaux. Dans ce papier, nous étudions une architecture simplifiée de réseaux dans laquelle lorsque deux routeurs sont reliés par un lien, les deux équipements extrémités de ce lien doivent être allumés. Chaque équipement ayant une consommation dépendant plutôt de son activation que de la quantité de trafic, notre objectif est de minimiser le nombre total d'équipements réseaux activés. Autrement dit, ce problème revient à effectuer un routage des demandes en minimisant le nombre d'arêtes dans la topologie. Nous proposons un programme linéaire pour résoudre ce problème et montrons des bornes simples sur des topologies particulières telles que la grille ou le graphe complet. Nous proposons ensuite une heuristique dont nous évaluons les performances à l'aide de simulations sur des topologies réelles et sur la grille.

Mots-clés : GMPLS, réseaux optiques, empilement de labels, hypergraphes, algorithmes d'approximation, programmation dynamique.

1 Introduction

With the recent increase of energy cost and the new sensibility of the public, governments and corporations towards energy consumption, the minimization of network energy consumption has become a priority. Some recent studies [CSB⁺08, MSBR09] exhibit that the traffic load of the routers only has a small influence on their energy consumption. Hence, the dominating factor is the number of switched-on network elements: interfaces, platforms, routers, . . . The goal thus is to find a routing that minimizes the (weighted) number of network elements used by the routing.

Looking at the costs, to minimize energy, we should try to use as few network elements as possible. Nevertheless, in most of the practical networks, PoPs or even routers can not be turned off. As a matter of fact, first, they are the source or destination of demands; second, they can be part of backup routes to protect the network against failures. Hence, for most practical cases, the goal will be to minimize the weighted number of network equipments.

In this paper, we consider a simplified architecture where a connection between two routers is represented as a link joining two network equipments. We can thus spare energy by turning off the two network equipments which are the extremities of a link connecting two routers. Therefore, the goal in this simplified architecture is *to find a subgraph minimum in number of links to route the demands*.

The contributions in this paper are the following:

- We model the problem of minimizing the routing energy consumption as a *linear program*.
- We give explicit close formulas or bounds for specific topologies, such as paths, complete graphs, and square grids. They provide limit behaviors and give indications on the application on practical networks.
- We present heuristics to find solutions for general networks. These heuristics are validated by comparing them with theoretical bounds.
- We study the *energy gain* on a set of topologies of existing backbone networks and show how the bounds provide closed values to the simulations.
- Finally, we discuss the impact of energy-efficient routing on *route length* and *fault tolerance*.

1.1 Related Work

Measure of energy consumptions. Several measurement campaigns of network energy consumption have been carried out in the last few years, see for example [MSBR08, MSBR09] and [CSB⁺08]. Their authors claim that the consumption of network devices is largely independent of their load. In particular, in [CSB⁺08], the authors were interested by routers' energy consumption. They observe that for the popular Cisco 12000 series the consumption

at a load of 75% is only 2% more than at an idle state (770W vs. 755W). In [MSBR09], the authors show through experimentations that the power consumed depends on the number of active ports. Explicitly disabling unused ports on a line card reduces the device power consumption.

Energy minimization. In [CSB⁺08], the authors model this problem as a linear program. The objective function is a weighted sum of the number of platforms and interfaces. They test on different networks to see how much energy can be saved with this model. However, they do not give intuitions nor explanations nor formulas for their results. In [IOR⁺10], the authors propose a rerouting at different layers in IP-over-WDM networks for energy savings while [PVC⁺09] study the impact of the technology for energy efficient routing. In [NIRW08, GS07, GS03], researchers proposed techniques such as putting idle subcomponents (line cards, ports etc.) to sleep, as well as adapting the rate for forwarding packets depending on the traffic in local area networks. In [CNR10], the authors propose a modulation of the radio configurations in fixed broadband wireless networks to reduce the power consumption.

In this paper, we first present the problem modeling as a linear program in Section 2. Then, specific topologies, such as trees or complete graphs, are discussed in Section 3. We show the good performance of our proposed heuristic in terms of power consumption in Section 4 and discussed of the impact of such algorithms for network operators in Section 5.

2 Problem Modeling

One way to reduce the power consumption induced by network equipments such as routers, consists in minimizing the number of turned-on network equipments. We model a network topology as an undirected graph $G = (V, E)$, with capacities c_e for edge $e \in E$ and we have to find a routing satisfying the set of demands D .

2.1 Linear Program

The presented problem can be modeled as a multicommodity *integral* flow problem in which the objective function is the minimization of the number of edges. For each edge $e \in E$, we introduce a binary variable x_e which says if the edge e is used or not: $x_e = 0$ if $f_e = 0$ and $x_e = 1$ if $f_e > 0$.

The objective function is then $\min \sum_{e \in E} x_e$, and the constraints are the following:
Flow constraints: $\forall (s, t) \in V \times V, \forall u \in V$,

$$\sum_{\forall v \in V^-(u)} f_{vu}^{st} - \sum_{\forall v \in V^+(u)} f_{uv}^{st} = \begin{cases} -\mathcal{D}_{st} & \text{if } u = s, \\ \mathcal{D}_{st} & \text{if } u = t, \\ 0 & \text{otherwise.} \end{cases}$$

Capacity constraints: $\forall (u, v) \in E, \sum_{d \in \mathcal{D}} (f_{uv}^d + f_{vu}^d) \leq x_e c_e$.

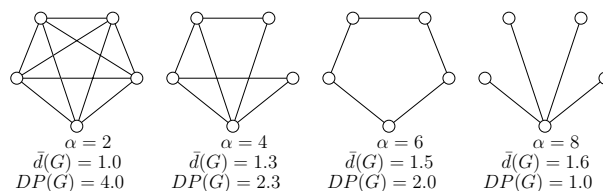


Figure 1: A Toy example, study of the complete graph with 5 vertices: subgraphs with the minimum number of edges, with α the capacity/demand ratio, $\bar{d}(G)$ average route length and $DP(G)$ the average number of distinct paths between two nodes.

2.2 Problem complexity.

Related optimization problems.

The problem presented in this paper is a special case of different well known network optimization problems: minimum routing, minimum concave flow problem [GP90], minimum flow problem with step cost functions. In operation research, this problem can be seen as a special case of what is named in the literature the FIXED CHARGE TRANSPORTATION PROBLEM [Spi64, Dan63], where the cost of the flow unit on an edge is zero. The problem is known to be NP-Complete, see for example [GP90]. Note that this problem is a particularly hard NP-problem: the number of possible subgraphs to test is strongly exponential for most graphs. Moreover, even for a given subgraph (when the set of edges to be used is fixed), the feasibility of a multicommodity integral flow problem has to be assessed and this problem is known to be NP-complete even for two commodities. Last, note that it is also the worst case of step-functions as most of the approximations by linearization are very far from a feasible solution.

In practice, we ran the solver CPLEX¹ to solve the linear program on different topologies. It takes a prohibitive time on large networks and only small to medium sized networks can be solved exactly. As an example, using an Intel Xeon 2.33GHz SMP X86_64 machine running under LINUX, we solved the 3-grid in 2 seconds in average, it took 9 hours to solve one instance of the problem on the 4×4 square grid (16 nodes, 24 edges). Larger grids could not be solved.

Hence the need of: First, deriving bounds on the solutions giving indications on the goodness of the solutions. Second, finding good heuristic algorithms giving near optimal routings.

3 Topology Study: extreme cases.

We present here the general framework of the studies of this paper. We then study two extreme cases for real networks, namely trees and complete graphs. They give us the limit behavior of the real networks for different network loads. We also derive an upper bound on the number of edges that can be spared.

¹<http://www-01.ibm.com/software/integration/optimization/cplex/>

3.1 Our Framework.

In practical networks, the demands vary during the life of the networks, e.g. with the increase of the number of users, with the development of new technologies, or, more simply, according to the time of the day. The goal of our study is to see how much energy can be spared depending on the resources available for the routing. In the framework we propose, each edge of the network has the same amount of capacity (edge capacity c) directly related to the volume of traffic exchanged between node (volume of demand d). This ratio between demand and capacity is expressed with the capacity/demand ratio α , where for all edges in the topology, $c = \alpha * d$, where c is the capacity of the edges and d is the volume of traffic exchanged between nodes. We perform an all-to-all routing with demands of volume d , that will stress our heuristic to its limit and also will keep the network connected. Then, for each topology, we look at the number of edges that can be spared for different capacity/demand ratios α . This ratio is linked in some contexts to the network *overprovisioning* factor.

As an illustration, Figure 1 shows an optimal solution given by the linear program for a complete graph with 5 nodes when the capacity/demand ratio α varies from 2 to 8. When α equals 2, all the links of the graph are needed to perform an all-to-all routing. The larger the ratio is, the fewer network equipments are needed until we reach the star graph where no more network equipments can be removed. The two extreme cases are represented as the complete graph ($\alpha = 2$) and for the tree ($\alpha = 8$).

Note that an orthogonal way to conduct the study is to fix the number of network equipments/edges to be turned-on and to compute the *load of the network*, that is the minimal capacity needed to be able to satisfy the all-to-all routing:

Definition 3.1 (Load of a Graph). *Let $G = (V, E)$ be a graph and $\mathcal{D} : V \times V \rightarrow \mathbb{R}$ a demand. The load of the network is the minimum over all routings (feasible flows \mathcal{F}) of the maximum load over all edges: $\min_{f \in \mathcal{F}} \max_{e \in E} f_e$.*

3.2 Load of the Minimal Subgraph: a Spanning Tree

In the undirected case, the subgraph with the minimum number of edges is a tree, as it is the smallest connected subgraph, see e.g. Figure 1. This minimal configuration can be attained when the capacity is larger than the load given in the next lemma:

Lemma 3.1. *[Tree and Spanning Tree]*

- *The load of the tree is $2v(n - v)$, where v is the size of the larger branch incident to the tree centroid.*
- *In a graph with n nodes and of maximum degree d , the load of a spanning tree is at least*

$$2 \left\lceil \frac{n}{d} \right\rceil \left(n - \left\lceil \frac{n}{d} \right\rceil \right).$$

Note that the load of a (spanning) tree mostly depends on the maximum node degree of the underlying graph. On a complete graph where $d = n - 1$, the tree with the lowest load is

a star and its load is $2(n-1)$, to be compared with the load of a path $\lfloor \frac{n^2}{2} \rfloor$, for which $d = 2$. Hence, networks with nodes of large degrees tend to attain the minimum configuration for smallest capacities, see Section 4.

3.3 Complete Graph - Bound on the Number of Spared Edges

In this subsection, we consider the complete graph with n nodes and $n(n-1)/2$ edges. This topology is the *second extreme case* (all possible edges) of our problem. It corresponds to the *design problem* of finding the best network satisfying the load when all possible edges between nodes can be used.

The all-to-all routing is possible on the complete graph as soon as the capacity/demand ratio α is larger than 2. We also have seen in the previous section that when the load is larger than $2(n-1)$ the routing is possible on the star with only $n-1$ edges, corresponding to the small fraction $2/n$ of the total number of edges of the complete graph. But what happens between these two capacities?

Lemma 3.2. *In a graph with n vertices, the all-to-all routing uses at least*

$$\max\left(\frac{2n(n-1)}{c+2}, n-1\right)$$

edges, with c the capacity of the edges.

Proof. When all the edges have the same capacity and for all-to-all demands, $kc \geq \sum_{i,j \in V \times V} d(i,j)$, where c is the link capacity, and $d(i,j)$ is the shortest path distance between i and j . Note that $\sum_{i,j \in V \times V} d(i,j)$ is given as a closed formula by twice the Wiener Index [Wie47].

In a complete graph, all the paths have length one. If we remove the edge ij , we know that two paths (from i to j and from j to i) are now of length at least 2. Hence, for a subgraph with k edges, we have at most $2k$ paths of length 1, and $n(n-1) - 2k$ paths of length at least two. This gives

$$kc \geq \sum_{ij \in V \times V} d(i,j) \geq 2k + 2(n(n-1) - 2k).$$

The bound, holding for any topology, gives $k \geq \frac{2n(n-1)}{c+2}$.

□

Validation. We validate in Figure 2 the results of the linear program given by CPLEX with the lower bound which is close to the optimal solution. The gain in terms of network equipments can reach 60%, indeed, only 4 edges (or 8 network equipments) are needed instead of 10 (or 20).

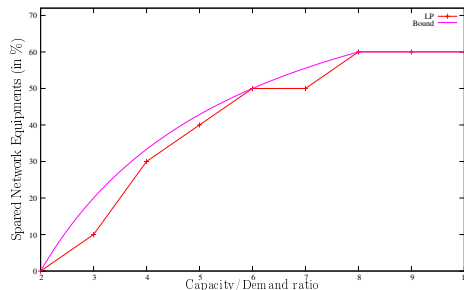


Figure 2: Number of spared edges on the complete graph with 5 nodes. Bound of Lemma 3.2 and linear program.

4 Results

After providing some bounds on extreme cases, such as trees or complete graphs, we propose in this section two heuristics for energy-efficient routing for classical real network topologies. In our experiments, we explore how many network equipments can be spared for different ranges of overprovisioning factor. We consider a range of capacities starting from the smallest capacity allowing to route all the network requests (overprovisioning factor OF equals 1) to the capacity allowing to route on a minimal subgraph, that is a spanning tree. We also study the impact of this routing on the route lengths and on the network fault tolerance.

We first consider the square grid which is an example of a structured network which often appears in practice for example in the context of wireless mesh networks. Then we study ten classical network topologies extracted from SNDLib website (<http://sndlib.zib.de>).



4.1 Heuristics

Two heuristics were run for the simulations. The Algorithm 1 presents a simple heuristic for our problem. We start from the whole network and we try to remove edges. The idea is to remove in priority edges that are less loaded. We believe that it is better to remove these edges that are not involved in many shortest paths than overloaded edges. The second one, called *Random* in plots, is used as a measure of comparison. It selects uniformly at random the links to be removed. Results of the *Random* heuristic are given averaged over 2000 runs.

4.2 Square Grid

We present results for the 4×4 -grid (LP, lower bound, construction and heuristics) and for a 10×10 -grid (lower bound, construction and heuristics), see Figure 3 and 4.

Limit Configurations. We first look at the load of the two limit configurations: the full grid and spanning trees. The load of the best spanning tree is given by the following lemma.


 Figure 3: Percentage of spared network equipments for a 4×4 grid.

 Figure 4: Percentage of spared network equipments for a 10×10 grid.

Algorithm 1 Heuristic by little loaded edge removal

Require: A graph $G = (V, E)$ with a capacity/demand ratio α ($\forall e \in E, c(e) = \alpha * d$), and a set of demands \mathcal{D} , each demand with a volume of traffic d .

while Edges can be removed **do**

Remove the edge e' that has not been chosen once, with the smallest value $\frac{c(e')}{r(e')}$, not disconnecting G .

Update the residual capacity $r(e), \forall e \in E$ while performing a routing of the demands

If $\exists e \in E$ such that $r(e) > c(e)$, then put back e' in G

end while

return the subgraph G .

Lemma 4.1 (Grid Spanning Tree). *The tree with the smallest load on a $a \times a$ square grid ($a \geq 3$) is a tree with a centroid of degree 4 and 4 branches of almost equal sizes. Its load is*

$$2 \left\lceil \frac{a^2 - 1}{4} \right\rceil \left(a^2 - \left\lceil \frac{a^2 - 1}{4} \right\rceil \right) \approx \frac{3a^4}{8}.$$

For example, for the 4×4 square grid (see Figure 3), the first capacity at which we attain the tree (and save the maximum number of edges) is 96. For this capacity, we use $a^2 - 1 = 15$ edges, when the full grid has 24 edges, saving 37.5% of edges.

Minimum Cut. On Figure 3, we have plotted the path length bound (magenta $-$). We see that, on the contrary to the complete graph, this bound is not very tight. The explanation is that edges in the square grid are not all equivalent: the edges in the middle of the grid have a larger load than the one on the borders with a shortest path routing.

The load can be computed by looking at the *minimum cut* of the network. If we consider the cut composed by all horizontal edges in the middle of the grid, the grid is divided into two sets with $a^2/2$ elements. The capacity of the cut has to be larger than the $2(a^2/2)^2$ demands. We get $c \geq a^3/2$. For the 4×4 grid, we can check that the load induced by the routing is $4^3/2 = 32$. This cut phenomena also explains that as soon as the routing is possible (for $\alpha = 32$), some edges can be spared (around 8%): the middle edges are fully loaded, but not the edges on the borders. The same cut phenomenas are present for some SNDLib topologies studied in the next section.

How many network equipments can be spared? For values of α larger than 32, a large number of edges can be saved at the beginning when the capacity/demand ratio decreases: 25% for an overprovisioning factor of 1.5 ($\alpha = 48$) and 33% for a factor of 2 ($\alpha = 64$). The savings become less important for larger capacities: only a 4 percent difference between capacities of 64 and 96, and the tree is attained with an overprovisioning factor of 2.9.

Behavior of the heuristics. Note that the heuristic (green \times) behaves well and is close to the optimal values given by the linear program (red $+$). The heuristic is also significantly better than the *Random* heuristic. For example, for a capacity/demand ratio of 40, the LP save 16%, the heuristic 13% and only 3% for *random*, when for a capacity of 55 and 70, we have respectively 25%, 22%, 18% and 33%, 33%, 26%.

4.3 SNDLib Topologies

Topologies. We studied here ten classical real network topologies extracted from the SNDLib website. These networks correspond to US (*Atlanta*), European (*Nobel EU*, *Cost266*), or single country (*Nobel Germany*, *France*) topologies. Their sizes span from 15 to 54 nodes and from 22 to 89 edges, as summarized in Table 1. For these 10 topologies, we computed energy efficient routings for different capacity/demand ratios. We could only run the linear program for the smallest network, *Atlanta*, for which the results are presented in Figure 5. (We see that the heuristic performs well on the topology, attaining the optimal value most of the time, when the *Random* heuristic needs a larger capacity to attain the same value.) As a matter of fact CPLEX already takes several hours on *Atlanta* to solve the problem for one capacity/demand ratio. We thus present the results found by the heuristic (which takes only tens of milliseconds) in Tables 1 and 2.

Limit Configurations. We first study the limit configurations. We report in Table 1 (b) the first capacity/demand ratio α_1 for which the heuristic can perform an all-to-all routing. As explained in Section 4.2, α_1 depends on the minimum cut of the network. We reported for




Figure 5: Percentage of spared network equipments for the *Atlanta* network.

each network a cut dividing the network into two almost equal-sized regions. The bound on α_1 implied by the cut is given in the second column of the table. We see that it gives a very good indication of α_1 for most of the networks, even if the value is not tight, as the heuristic does not always find the optimal solution. For example, we see that the bound is tight for *Atlanta*, where the minimum cut is of size 3 and splits the network in two sub-networks of sizes 8 and 7. For *Atlanta*, the demand/capacity ratios evaluated by the minimum cut bound is $\alpha_1 = 38$ (compared to 38) and *Nobel EU*, where we get $\alpha_1 = 131$ (compared to 131).

As explained in Section 3 (Lemma 3.1), α_{tree} , for which we get a spanning tree, depends on the network node degree. On the contrary to regular network, such as the grid, the existence of a spanning tree centered on the node of max degree with equal-sized branch is not given. Hence, the bounds given in Table 1 (b) are not tight. Nevertheless, we see that the node degree still is a good indication on what can be achieved on these networks. Graphs with low maximum degree attains the best configuration for lower values of capacity/demand ratio. For example, the network *Atlanta* has a maximum node degree of 4 (the bound is 88) and α_{tree} is 108, when for *Zib54* α_{tree} is 1386 for a maximum degree of 10.

Results: Spared network equipments. We give the percentage of spared network equipments in function of the overprovisioning factor. A factor of one means that we use the minimal capacity necessary to route all the demands (this capacity corresponds to the value α_1), when for example a factor of two means that we have twice this minimal capacity. Note that in most today's backbone networks overprovisioning is heavily used as it is an efficient and easy way to provide protection against failure: links are often used between 30 and 50 % of their capacity.

First, note that on some of the ten topologies, as soon as the routing is feasible (overprovisioning equals 1), some network equipments can already be turned off (12% for Norway and Nobel EU). As a matter of fact, in this case, the important edges of the network (the edges in the minimum cut for example) are fully used, but at the same time edges at the periphery are less used and some can be spared. With an overprovisioning factor of two, around one third of the edges can be spared (and even 53% for the Pioro40 network). With larger factors (3 or 4), the gain is not as important, but still some network equipments can

			Overprovisioning factor				Tree	
			1	2	3	4	OF	%SNE
Atlanta	15	22	0%	32%	36%	36%	2.66	36%
France	25	45	0%	42%	44%	47%	3.13	47%
Nobel EU	28	41	12%	32%	34%	34%	2.76	34%
Nobel Germany	17	26	0%	35%	39%	39%	2.75	39%
Cost266	37	57	3.5%	32%	35%	37%	3.65	37%
Guil39	39	86	0%	45%	50%	52%	8.25	56%
Norway	27	51	12%	43%	47%	47%	4.71	49%
Zib54	54	80	0%	30%	33%	33%	4.71	34%
Pioro40	40	89	0%	53%	54%	55%	5.12	56%
New York	16	49	2.0%	59%	63%	67%	5.2	69%

(a) Gain of network equipments (in %).

Simulations		Values given by the bounds			
α_1	α_{tree}	α_{tree}		α_1	
		Cut	Bound	d_{max}	Bound
38	108	3 (8/7)	38	4	88
67	210	5 (10/15)	60	10	132
131	384	3 (14/14)	131	5	264
44	123	4 (9/8)	36	5	104
175	651	4 (19/18)	171	4	540
85	745	12 (20/19)	64	8	340
75	370	5 (12/15)	72	6	220
294	1386	10 (27/27)	146	10	576
153	783	8 (18/22)	99	5	512
15	78	11 (7/9)	12	11	56

(b) Evaluation of the load given by the bounds.

Table 1: (a) Gain of network equipments (in %) depending on the overprovisioning factor and (b) Evaluation of the load given by the bounds in previous section with α_1 the capacity/demand ratio for overprovisioning factor equals 1 and α_{tree} for the tree.

	V	E	Overprovisioning factor				Tree	
			1	2	3	4	OF	SF
Atlanta	15	22	1.000	1.186	1.255	1.255	2.66	1.255
France	25	45	1.000	1.096	1.123	1.167	3.13	1.158
Nobel EU	28	41	1.080	1.143	1.241	1.242	2.76	1.247
Nobel Germany	17	26	1.000	1.112	1.176	1.132	2.75	1.179
Cost266	37	57	1.045	1.112	1.195	1.314	3.65	1.319
Guil39	39	86	1.000	1.179	1.208	1.280	8.25	1.497
Norway	27	51	1.024	1.170	1.184	1.205	4.71	1.248
Zib54	54	80	1.000	1.019	1.067	1.097	4.71	1.110
Pioro40	40	89	1.000	1.254	1.320	1.382	5.12	1.424
New York	16	49	1.005	1.240	1.262	1.282	5.2	1.325

(a) Average multiplicative stretch factor

	Overprovisioning factor				Tree	
	1	2	3	4	OF	DP
	2.352	1.095	1.000	1.000	2.66	1.000
	2.483	1.023	1.010	1.000	3.13	1.000
	1.817	1.074	1.000	1.000	2.76	1.000
	2.353	1.044	1.000	1.000	2.75	1.000
	2.467	1.122	1.068	1.000	3.65	1.000
	3.677	1.412	1.138	1.065	8.25	1.000
	2.613	1.137	1.043	1.009	4.71	1.000
	2.164	1.048	1.015	1.002	4.71	1.000
	4.058	1.124	1.086	1.027	5.12	1.000
	4.892	1.242	1.190	1.025	5.2	1.000

(b) Average number of disjoint paths

Table 2: Impact of the energy-efficient routing on (a) the route length and on (b) the network fault tolerance.

be saved (e.g. 36% for Atlanta). We show in Table 2(a), in the 2 last columns, the overprovisioning factor (OF) for which the tree is attained together with the corresponding spared network equipments in percentage (SNE). The values are directly linked to the density of the network. For example, *Norway* needs a factor of 4.71 to reach the tree with 26 links (instead of 51), sparing 49% edges. Hence, the larger the density of the network, the more the network equipments that can be turned off.

To conclude, for all the studied networks *between one third and one half of the network equipments can be spared for usual overprovisioning factors*. Furthermore, when the best routing cannot be found by the linear program (large topologies), the *proposed heuristic found close to optimal solutions*.

4.4 Impact on the Network

We believe that network operators will implement energy efficient routing only if the impact on other parameters is limited. We discuss here of the impact on the route lengths and on the fault tolerance of our proposed heuristic.

Impact on the route lengths. When turning-off some components in a network, we save some energy but, at the same time, we route on longer paths. The *multiplicative stretch* is defined as the ratio between the average route length in the new routing divided by the average route length with the old routing (using all the edges). Results for the SNDLib topologies are given in Table 2 (a). A stretch of one correspond to the cases where no edge could be spared and thus the routing is not affected. We see that, as expected, the general trend is that when the overprovisioning factor increases, the paths become longer.

Nevertheless we see that the impact on the route lengths is limited. For example, for the topology *Zib54*, the increase is less than 10% for the extreme case when routing on a tree (for 33% of turned-off network equipments). In general, the increase for this extreme case spans from 10% to 31%, but most of the time around 20%. For the network with the larger impact, *Pioro40*, the stretch is increased of 38% for the tree, with a saving of 55% of the edges. We see that a saving of already 53% of the edges is attained for an overprovisioning factor of 2, leading to a route length increase of only 25%.

Impact on the fault tolerance. We measure the network fault tolerance as the *average number of disjoint paths* linking two nodes. Results are shown in Table 2 (b). We see that the full network topologies have a number of disjoint paths between two and four. When routing on the tree, the number is of course one as only one route exists. The drop is quick as all the ten networks have a number below two for an overprovisioning factor of 2.

Discussion on Technology. When there is a failure between two nodes with only one route, it may be necessary to turn on some network equipments to find a new route. Hence, the use of such energy efficient solutions is conditioned by the existence of *technologies allowing a quick switching on of network equipment* and network links. Network equipments companies are currently working on designing this kind of equipments [Pap08].

5 Summary of the simulations, conclusion and perspectives

We presented through a simplified architecture the problem of minimizing power consumption in networks. The simulations on real topologies show that the gain in energy is significant when some network equipments can be turned-off.

- At least one third of the network equipments can be spared for usual range of demands.
- If we use the values detailed [CSB⁺08] for the power consumption of the network equipments (approximately 800 Watts), the reduction in terms of number of Watts for *Cost266* which has 37% of spared network equipments is *33kWh*. This leads to a reduction of power consumption of approximately *300MWh* per year.
- The *route lengths* increase, but not too much: around 20% for almost all studied topologies.
- *Fault tolerance* can be achieved with the use of *fast switching-on technologies* or by adding *disjoint path constraints* to the problem implying for example a *hamiltonian cycle* as a better configuration than a tree.

As part of future work, we plan to study a detailed cost function for a more complex router architecture. Moreover, we will carry on lab experimentations on small network topologies to measure in practice the performance of the proposed energy-efficient routing.

Acknowledgements

This work has been partly funded by the ANR DIMAGREEN².

References

- [CNR10] D. Coudert, N. Nepomuceno, and H. Rivano. Power-efficient radio configuration in fixed broadband wireless networks. *Computer Communications*, 2010. to appear.
- [CSB⁺08] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright. Power awareness in network design and routing. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 457–465, 2008.
- [Dan63] G.B. Dantzig. *Linear programming and extensions*. Princeton Univ Pr, 1963.

²<http://www.sop.inria.fr/teams/mascotte/Contrats/DIMAGREEN/wiki/>

- [GP90] G.M. Guisewite and P.M. Pardalos. Minimum concave-cost network flow problems: Applications, complexity, and algorithms. *Annals of Operations Research*, 25(1):75–99, 1990.
- [GS03] M. Gupta and S. Singh. Greening of the internet. In *ACM SIGCOMM*, 2003.
- [GS07] M. Gupta and S. Singh. Energy conservation with low power modes in Ethernet LAN environments. In *IEEE INFOCOM (Mini Symposium)*, 2007.
- [IOR⁺10] F. Idzikowski, S. Orłowski, C. Raack, H. Woesner, and A. Wolisz. Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios. In *ONDM*, 2010.
- [MSBR08] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan. Energy Aware Network Operations. In *HP Labs*, 2008.
- [MSBR09] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan. A Power Benchmarking Framework for Network Devices. In *IFIP NETWORKING 2009*, pages 795–808, may 2009.
- [NIRW08] L. Nedeveschi, S. and Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall. Reducing Network Energy Consumption via Rate-Adaptation and Sleeping. In *NSDI*, 2008.
- [Pap08] White Paper. Energy efficiency for network equipment: Two steps beyond greenwashing. Technical report, Juniper Networks Inc., 2008.
- [PVC⁺09] B. Puype, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester. Power Reduction Techniques in Multilayer Traffic Engineering. In *ICTON*, 2009.
- [Spi64] Kurt Spielberg. On the fixed charge transportation problem. In *Proceedings of the 1964 19th ACM national conference*, pages 11.101–11.1013, New York, NY, USA, 1964. ACM.
- [Wie47] H. Wiener. Structural determination of paraffin boiling points. *Journal of the American Chemical Society*, 69(1):17–20, 1947.



Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399