

Robust Redundancy Elimination for Energy-aware Routing

David Coudert
Joint Project COATI
I3S (CNRS - UNS)

INRIA Sophia Antipolis, France
david.coudert@inria.fr

Arie M. C. A. Koster
Lehrstuhl II für Mathematik

RWTH Aachen, Germany
koster@math2.rwth-aachen.de

Truong Khoa Phan
Joint Project COATI
I3S (CNRS - UNS)

INRIA Sophia Antipolis, France
truong_khoa.phan@inria.fr

Martin Tieves
Lehrstuhl II für Mathematik

RWTH Aachen, Germany
tieves@math2.rwth-aachen.de

Abstract—Many studies have shown that energy-aware routing (EAR) can significantly reduce energy consumption of a backbone network. Redundancy Elimination (RE) techniques provide a complementary approach to reduce the amount of traffic in the network. In particular, the GreenRE model combines both techniques, offering potentially significant energy savings.

We propose a concept for respecting uncertain rates of redundant traffic within the GreenRE model, closing the gap between theoretical modeling and drawn-from life data. To model redundancy rate uncertainty, the robust optimization approach of Bertsimas and Sim (2004) is adapted and the problem is formally defined as mixed integer linear program.

An exemplary evaluation of this concept with real-life traffic traces and estimated fluctuations of data redundancy shows that this closer-to-reality model potentially offers significant energy savings in comparison to GreenRE and EAR.

Keywords: Robust Network Optimization, Green Networking, Energy-aware Routing, Redundancy Elimination.

I. INTRODUCTION

In modern communication infrastructures, energy-consumption is one of the most critical aspects for designing network topologies. Routing has not only to be feasible with respect to congestion, but as energy efficient as possible. Therefore, the classical energy-aware routing (EAR) problem aims at minimizing the active elements of routers (the most influencing factor of energy consumption), while all traffic demands are routed without any overloaded links [10, 11, 16, 26]. Recently, the GreenRE model [14] has been proposed, combining EAR and Redundancy Elimination (RE) techniques to increase energy efficiency of a backbone network.

Although solving the GreenRE model is already a complex task [20], it does not take traffic redundancy fluctuations into account. Instead, each of the demands contains a constant factor of redundant traffic. This assumption may lead to infeasible or inefficient network designs, i.e., a high value of estimated traffic redundancy causes overloading, whereas using an underestimated value wastes energy savings.

The contribution of this paper is an extension of the GreenRE model as state-of-the-art technique to include uncertainty of traffic redundancy as well. Therefore, a mean of dealing with uncertainties has to be chosen carefully. While a general worst-case analysis is inefficient in applications, the

Γ -robustness concept [7] models uncertainties in a more realistic way. This technology-independent concept has already been successfully applied to, for example the network design problem under demand uncertainty [1, 18]. Given a parameter $\Gamma \geq 0$, the problem considers any simultaneous deviation of at most Γ traffic pairs from their nominal traffic volumes.

In this paper, we extend the GreenRE model by applying the idea of Γ -robustness to uncertain data redundancy. We propose GreenRobustRE - a model that includes uncertainty of redundancy elimination rates. Accordingly, contributions are structured as:

- In Section III, we define and formulate the GreenRobustRE problem as mixed integer linear program. To the best of our knowledge, this is the first work considering robustness on redundancy elimination for traffic flows.
- In Section IV, we exemplarily evaluate energy savings for two networks based on real-life traffic traces and estimated redundancy fluctuation. The results show a significant increase of energy savings by the GreenRobustRE model, compared to previous models.

As central point of this paper, we show the superiority of the GreenRobustRE model in both, being closer to reality (model wise - Section III) and yielding better solutions (Section IV). A representation as mixed integer linear program offers an accurate description of the potential of the proposed concept. For large networks, a more refined solution approach is necessary.

We start with a review of already known, related work (especially the GreenRE model). Therefore, we repeat the complementary concepts of energy aware routing (II.A) and redundancy elimination (II.B) before presenting the GreenRE as a combination of both ideas. Furthermore, we introduce the Γ -robust optimization approach as the background of the GreenRobustRE problem. Following we will explain the GreenRobustRE problem, concluding with exemplary computations and a conclusion/evaluation.

II. BACKGROUND: AN EVOLUTION OF MODELS

A. Classical Energy-aware Routing (EAR) Model

Many studies in literature have shown that energy consumption of a backbone network is mainly depending on

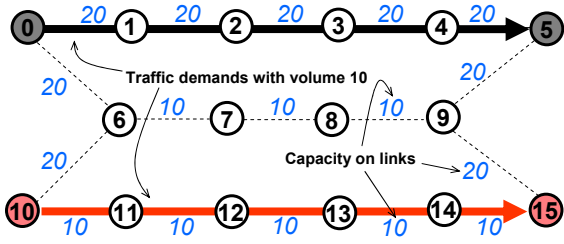


Fig. 1. Shortest path routing: turn off 7 links.

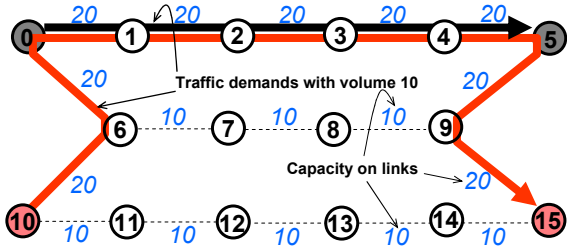


Fig. 2. Energy-aware routing: turn off 8 links.

the number of active elements of routers, such as ports, line cards, base chassis, etc. [10, 16, 21]. We consider a link as a connection between two network interfaces, one in each end-router. Therefore, the classical EAR was proposed to aggregate traffic flows into fewer links while preserving connectivity and QoS. Then, links that do not carry any traffic will be turned off (or more precisely, two network interfaces on the two routers will be turned off) to save energy for the network. A visualization is given in Fig. 1 and Fig. 2. Assume two traffic demands (from router 0 to 5 and from 10 to 15), both with a volume of 10 Gbps. The shortest path routing (as the standard routing principle), is shown in Fig. 1. It constitutes 10 active links whereas the remaining 7 links can be turned off. However, taking energy consumption into account, in Fig. 2, the EAR solution allows to turn off 8 links, 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 [13], 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 [11, 13, 26].

B. Redundancy Elimination (RE)

As a complementary approach to energy-efficient via routing principles, many studies focus on reducing traffic load for the Internet [2, 3, 4, 24, 25]. Observing that some data is repeatedly requested between routers of a backbone network, a large amount of traffic contains redundancy. For example, popular contents, such as new movies, are often downloaded several times subsequently. As a result, traffic on the backbone contains large amounts of redundancy i.e., traffic that could be avoided by appropriate memory management [2, 4, 24].

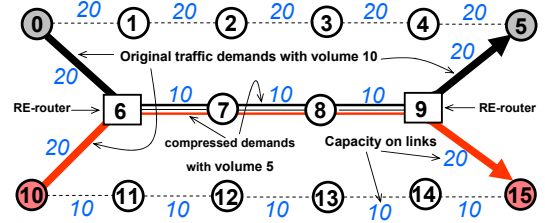


Fig. 3. GreenRE routing with 50% of traffic redundancy: turn off 10 links, enable 2 RE-routers.

Removing this redundancy virtually increases the capacity of network links allowing for an efficient EAR.

Therefore, the protocol independent redundancy elimination technique has been proposed: data traffic is splitted into small chunks and each chunk is associated with a small key. To reduce link loads, routers may cache recent data chunks and replace repeated parts on incoming traffic by small keys. Forwarding these keys only, the original data will be recovered further on, given a cache synchronization between the sending and the receiving routers. Anand et al. [2, 3] and Song et al. [24] have shown that using real traffic traces, RE can reduce link loads by 10% - 50%. This technique has not been deployed on backbone networks yet. However, the industry has adopted RE in devices called Wide-Area Network Optimization Controller [9, 15, 23]. In the following, a router with a RE-device is called RE-router. It can perform as both, encoder: compression of outgoing (to another node) and decoder: de-compression of incoming (from another node) traffic, if requested to do so. We will shortly option this RE-service.

C. GreenRE Model

The GreenRE model is an extension of EAR, i.e., a combination of RE and EAR. In this model, the redundancy elimination technique virtually increases the capacity of the network. A drawback is, that the caching process increases the energy consumption. The authors in [14] have shown, that a router performing RE consumes more energy than usual. This introduces a trade-off between enabling RE on routers (increasing their power consumption) and turning off links (saving their expenses), such that designing an optimal network topology is not trivial.

As proof of concept, we refer to Fig. 3. Let a RE-router cost 30 Watts [14] and a link consume 200 Watts [11]. Assume that 50% of the traffic is redundant and RE-service is enabled at router 6 and router 9. Hence, all traffic flows passing between the routers 6, 7, 8, 9 can be compressed to 5 Gbps at router 6 and are de-compressed to full size at router 9. So, the routing as shown in Fig. 3 is feasible (without any congestion). As a result, the GreenRE solution allows to turn off 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. 2).

More precisely, the GreenRE problem is defined on an undirected graph $G = (V, E)$, where C_e denotes the ca-

capacity of link $e \in E$. The set of demands is given by $D = \{(s, t) \in V \times V : s \neq t\}$ and $\mathcal{D}^{st} \geq 0$ denotes the amount of traffic requested from target t of source s . Let $PE_e, PN_u \geq 0$ be the power consumption of an active link / RE-router. The constant $\lambda^{st} \in [0, 1)$ denotes the percentage of traffic redundancy of a demand (s, t) . Corresponding to λ^{st} , we define $\gamma^{st} := (1 - \lambda^{st})$, which represents the percentage of unique (non redundant) traffic. For instance, for a 10 Gbps traffic demand with $\lambda^{st} = 40\%$ of redundancy, its volume can be reduced by GreenRE to $10\gamma^{st} = 6$ Gbps of non-redundant traffic. For simplicity, a traffic flow, from which redundancy has been removed, is called a *compressed flow*.

Binary variables x_{uv} and w_u denote the activated links / RE-routers. We use variables $f_{uv}^{st}, g_{uv}^{st} \geq 0 \forall (s, t) \in D, uv \in E$ describing the fraction of normal and compressed flows of demand (s, t) , routed directly from u to v . Denoting $N(u)$ for the set of neighbors of u , we formulate the GreenRE model as follows:

$$\min \sum_{e \in E} PE_e x_e + \sum_{u \in V} PN_u w_u \quad (1)$$

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

$$\sum_{(s,t) \in D} \mathcal{D}^{st} (f_e^{st} + \gamma^{st} g_e^{st}) \leq \mu C_e x_e \quad \forall e \in E \quad (3)$$

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

$$\sum_{v \in N(u)} (g_{vu}^{st} - g_{uv}^{st}) \leq w_u \quad \forall u \in V, (s, t) \in D \quad (5)$$

$$x_e \in \{0, 1\}, w_u \in \{0, 1\}, f_{uv}^{st} \in [0, 1], g_{uv}^{st} \in [0, 1] \quad (6)$$

where $f_e^{st} = f_{uv}^{st} + f_{vu}^{st}$ and $g_e^{st} = g_{uv}^{st} + g_{vu}^{st}$. The objective function (1) minimizes the power consumption of the network represented by the number of active links and RE-routers. The equations (2) establish flow conservation, whereas the constraints (3) limit the available capacity (where μ denotes the maximum link utilization). The constraints (4) and (5) determine, whether decoding/encoding is necessary at a node u , such that RE-service is activated ($w_u = 1$) or not. I.e., compression is necessary/takes place when the sum of incoming compressed flow is bigger than the sum of outgoing compressed flow (4) or vice versa (5). This necessity is recognized by a big M ($M = 1$) constraint. So, if u is a normal router, it only forwards flows without compression or de-compression and if the percentage of compressed flow changes in a node, a RE-router is required. For the sake of notation, we assume that, all routers have the capability to perform RE-service, so we can enable it when needed.

It is also noted, that in a feasible solution of GreenRE, a compressed flow is decompressed somewhere on the way to

its destination. Otherwise, one node (latest at the target) would receive more incoming compressed traffic as outgoing (without being a RE-router), violating constraints (5). Consequently, in every optimal solution, there will be at least two active RE-routers or none at all. Clearly, employing more RE-routers (or links) than absolutely needed is feasible but not optimal.

In an aggregated perspective, the above described models are a range of more and more fine-tuned concepts to model energy efficient networks. Automatically, this leads to questions related to quality measures of these models, which again is dependent on precise data. Since in most cases, data is uncertain by nature, we believe that this uncertainty has to be included within these models. Our contribution is a proposal of including uncertainties within the GreenRE model as state-of-the-art concept. However, before going into detail, we will shortly outline the Γ -robustness concept as chosen mean to deal with such uncertainties.

D. Robust Optimization

Over the past years, robust optimization has been established as a special branch of mathematical optimization allowing to handle uncertain data [5, 6]. A specialization of robust optimization, which is particularly attractive by its computational tractability, is the so-called Γ -robustness concept introduced by Bertsimas and Sim [7, 8]. Instead of deterministic coefficients, the coefficients a_j of a constraint $\sum_j a_j x_j \leq b$ are assumed to be random variables. Bertsimas and Sim have shown that in case all random variables are independent and have a symmetric distribution of the form $a_j \in [\bar{a}_j - \hat{a}_j, \bar{a}_j + \hat{a}_j]$ (with \bar{a}_j the average and \hat{a}_j the maximum deviation), it can be guaranteed that the constraint is satisfied with high probability by defining an appropriate integer Γ and replacing the constraint by

$$\sum_j \bar{a}_j x_j + \max_{J: |J| \leq \Gamma} \sum_{j \in J} \hat{a}_j x_j \leq b. \quad (7)$$

This constraint models that for each realization of the uncertainties at most Γ many (but arbitrary) coefficients can deviate from their nominal value. Given an arbitrary realization, it is shown in [7, 8], that the probability that (7) is violated, is about $1 - \Phi(\frac{\Gamma-1}{\sqrt{n}})$, where Φ is the cumulative distribution function of the standard normal distribution and n equals the number of uncertain coefficients. This result is independent of the actual distribution of a_j .

Note, that constraint 7 is deterministic and the complete problem can be reformulated as a standard mixed integer problem. So the model including uncertainty can be solved by the same means as the original problem, again see [7, 8] for details. From a practical perspective, by varying the parameter Γ , different solutions can be obtained with different levels of robustness (the higher Γ the more robust, but also more expensive, the solution is). This concept has already been applied to several network optimization problems (cf. [1, 12, 18]).

III. GREENROBUSTRE MODEL

As state-of-the-art model for energy-minimal routing, the deterministic GreenRE model assumes that each traffic de-

mand has a constant non-redundant value γ^{st} . This assumption leads to an inaccurate evaluation of energy savings, since the actual traffic redundancy rate fluctuates and is not known in advance. In practice, avoiding congestion is the most pressing matter, such that modeling has to be very close to worst-case analysis. By the above mentioned Γ -robustness and its probability bound, the conservatism of modeling can be alleviated by employing this concept. If the Γ is chosen appropriate, the probability of feasibility is high enough and as we show in Section IV, a significant improvement over the worst-case solution is still possible.

In the following, we propose the GreenRobustRE model, which addresses fluctuations by optimizing against a certain amount of uncertainties. As a consequence, the link capacity constraint (3) is deterministically satisfied, if this amount of uncertainty is realized (and satisfied by a very high probability for any other realization). Therefore, we adapt the approach of Bertsimas and Sim [7, 8] as follows: For each demand pair, two values describe the potential (or “realized” in the sense of random variables) redundancy elimination: a (nominal) default value $\bar{\gamma}^{st} \in (0, 1]$ and a maximal deviation $\hat{\gamma}^{st} \geq 0$, ($\bar{\gamma}^{st} + \hat{\gamma}^{st} \leq 1$), such that the actual redundancy value γ^{st} is known to be within $[\bar{\gamma}^{st}, \bar{\gamma}^{st} + \hat{\gamma}^{st}]$. So, whereas γ^{st} is a deterministic value in the GreenRE, it is now a random variable, symmetric distributed on an interval and defined by the two values $\bar{\gamma}^{st}$ and $\hat{\gamma}^{st}$. Potentially, each demand can be compressed by its default ratio to $\bar{\gamma}^{st} \mathcal{D}^{st}$. Applying Γ -robustness, we consider that at most Γ redundancy ratios fluctuate simultaneously. This means, the affected demand volumes have a lower compression potential, i.e., a higher value of γ^{st} . Consequentially, in Γ many cases the compressed flow can amount to a value as high as $(\bar{\gamma}^{st} + \hat{\gamma}^{st}) \mathcal{D}^{st}$.

For instance, based on historical traces, a demand (s, t) seems to contain 60% of non-redundant (unique) traffic on average. Hence, we assume a nominal non-redundant ratio of $\bar{\gamma}^{st} = 0.6$. Assuming at most 90% of the traffic at any time is non-redundant as an upper bound, we can protect ourself against wrong assumptions by adding $\hat{\gamma}^{st} = 0.3$. Depending on the desired level of protection of our solution, we choose a Γ -value, such that our solution is still feasible (and optimal) if at most Γ many redundancy ratios deviate their assumptions, without specifying which ones.

Given a parameter $0 \leq \Gamma \leq |D|$, the GreenRobustRE problem is to find a feasible routing at minimal energy costs, while the link capacity constraints are satisfied if at most Γ traffic pairs deviate from their $\bar{\gamma}^{st}$ values simultaneously. Note that $\Gamma = |D|$ amounts to worst-case optimization, whereas $\Gamma = 0$ models the opportunistic case without uncertainty. The straightforward (but nonlinear) robust capacity constraint for a given Γ and an edge $e \in E$ is:

$$\sum_{(s,t) \in D} \mathcal{D}^{st} (f_e^{st} + \bar{\gamma}^{st} g_e^{st}) + \max_{\substack{Q \subseteq D \\ |Q| \leq \Gamma}} \left\{ \sum_{(s,t) \in Q} \hat{\gamma}^{st} \mathcal{D}^{st} g_e^{st} \right\} \leq \mu C_e x_e \quad \forall e \in E \quad (8)$$

Given g_e^{st} , the maximum part of (8) can be computed by:

$$\begin{aligned} \beta(g, \Gamma) := \max \quad & \sum_{(s,t) \in D} \hat{\gamma}^{st} \mathcal{D}^{st} g_e^{st} \\ \text{s.t.} \quad & \sum_{(s,t) \in D} z_e^{st} \leq \Gamma \quad [\pi_e] \\ & z_e^{st} \in \{0, 1\} \quad [\rho_e^{st}] \end{aligned}$$

Based on [7], a *compact* reformulation can be obtained by employing total-unimodularity and LP duality of $\beta(g, \Gamma)$:

$$\begin{aligned} \beta(g, \Gamma) = \min \quad & \Gamma \pi_e + \sum_{(s,t) \in D} \rho_e^{st} \\ \text{s.t.} \quad & \pi_e + \rho_e^{st} \geq \hat{\gamma}^{st} \mathcal{D}^{st} g_e^{st} \quad \forall (s, t) \in D \\ & \rho_e^{st}, \pi_e \geq 0 \quad \forall (s, t) \in D \end{aligned}$$

where the primal binary variables z_e^{st} denote whether or not g_e^{st} is part of the subset $Q \subseteq D$. The dual variables π_e and ρ_e^{st} corresponds to the constraint $\sum_{(s,t) \in D} z_e^{st} \leq \Gamma$ and $z_e^{st} \leq 1$ (in the linear relaxation), respectively. Embedding this into (1)–(6), the GreenRobustRE can be compactly formulated by replacing the constraint (8) by:

$$\begin{aligned} \sum_{(s,t) \in D} (\mathcal{D}^{st} (f_e^{st} + \bar{\gamma}^{st} g_e^{st}) + \rho_e^{st}) + \Gamma \pi_e &\leq \mu x_e C_e \quad \forall e \in E \\ \pi_e + \rho_e^{st} &\geq \hat{\gamma}^{st} \mathcal{D}^{st} g_e^{st} \quad \forall (s, t) \in D, \forall e \in E \\ \rho_e^{st}, \pi_e &\geq 0 \quad \forall (s, t) \in D, \forall e \in E \end{aligned}$$

Compared to the deterministic model GreenRE which has $|E| + |V| + 4|E||D|$ variables and $|E| + 3|V||D|$ constraints, this GreenRobustRE model has $|E| + |E||D|$ additional variables and $|E||D|$ additional constraints.

Note, that by the above reformulation, we can obtain a new (deterministic) mixed integer problem (called GreenRobustRE), protecting against uncertainties with high probability. While we believe that the theoretical improvement of this model is apparent by the above explanations, we will give a computational/practical evaluation in the next section.

IV. COMPUTATIONAL EVALUATION

A. Test instances and Experimental settings

We solve the GreenRobustRE model with IBM ILOG CPLEX 12.4 solver [17]. All computations were carried out on a 2.7 Ghz Intel Core i7 with 8 GB RAM. We consider real-life traffic traces collected from the SNDlib [22]: the U.S. Internet2 Network (Abilene) ($|V| = 12$, $|E| = 15$, $|D| = 130$) and the national research backbone network Germany17 ($|V| = 17$, $|E| = 26$, $|D| = 251$). Capacity is set to $C = 10$ Gbps and $\mu = 0.5$ [11] for each link.

In our computations, we use a single traffic matrix consisting of the mean volume for each traffic demand during a one day period. To achieve a network with high link utilization, all traffic was scaled with a factor four, while to avoid individual bottlenecks, we use four parallel links for (Köln, Frankfurt) in the Germany17 network and use double links for four links in the Abilene network: (ATLang, HSTNng), (ATLang,

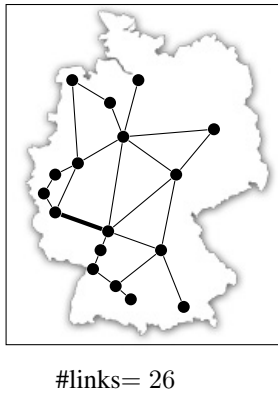


Fig. 5. EAR Solution Germany17

WASHng), (CHINng, IPLSng) and (HSTNng, LOSAng). For each network, 9 scenarios are generated by combining three nominal values $\bar{\gamma}$ and three deviation values $\hat{\gamma}$. In our tests, we assume that in every scenario the ranges of the compression values are independent of the node-pair, i.e., $\bar{\gamma}^{st} = \bar{\gamma}$ and $\hat{\gamma}^{st} = \hat{\gamma}$ for all $s, t \in D$. According to [2, 3] an upper bound on traffic redundancy of 50% can be assumed. Therefore, we assume that $\bar{\gamma} \geq 0.5$. In fact, we use three scenarios ($\bar{\gamma} = 0.5, \hat{\gamma} = 0.1$), ($\bar{\gamma} = 0.5, \hat{\gamma} = 0.25$) and ($\bar{\gamma} = 0.5, \hat{\gamma} = 0.5$) to represent traffic demands with high redundant ratio ($\bar{\gamma} = 0.5$) and low ($\hat{\gamma} = 0.1$), medium ($\hat{\gamma} = 0.25$) or high ($\hat{\gamma} = 0.5$) deviation. Similarly, the other scenarios are ($\bar{\gamma} = 0.7, \hat{\gamma} = 0.1$), ($\bar{\gamma} = 0.7, \hat{\gamma} = 0.2$), ($\bar{\gamma} = 0.7, \hat{\gamma} = 0.3$) and ($\bar{\gamma} = 0.8, \hat{\gamma} = 0.05$), ($\bar{\gamma} = 0.8, \hat{\gamma} = 0.1$), ($\bar{\gamma} = 0.8, \hat{\gamma} = 0.2$). For each scenario, we vary the robustness parameter Γ between 0 and $|D|$.

B. Results & Discussion

Before discussing particular trends or characteristics of solutions, we want to give a visualization of a typical solution of GreenRobustRE. In Fig. 4, we present solutions (within 10% optimality gap) for the Germany17 instance with $\bar{\gamma} = 0.7$ and $\hat{\gamma} = 0.2$ ($\Gamma \in \{0, 10, 15, 251\}$). The figure indicates by line thickness, that the edge Koeln-Frankfurt is always employed multiple times (3, 4, 4, 4). It is noted, that the $\Gamma = 0$ case mirrors the GreenRE model ($\gamma = 0.7$) and the $\Gamma = 251$ case equals to the GreenRE model with $\gamma = 0.9$. As above, $\gamma^{st} = \gamma$ for all demands $s, t \in D$. The subset of chosen edges is printed black and the activated RE-routers are displayed as red squares. For comparisons sake, Fig. 5 presents the corresponding EAR solution, i.e., routing without any compression/decompression (note that the edge between Frankfurt and Koeln has to be used 4 times).

1) *Energy savings vs. robustness*: In this section, we investigate the relation between energy savings and the level of robustness regarding the parameter Γ . All instances of the Abilene network can be solved to optimality in less than 10 minutes. For the Germany17 network, we limit the solving time to one hour and all best solutions are within 10% of optimality.

In a typical solution between three and seven RE-routers are activated. We observed that this number can change independently of the Γ value. A prognosis is difficult to give, since the number of RE-routers is highly dependent on the traffic volumes, the capacity, and the network topology. Clearly, the same holds for the employed edges and depending on the demands and the employed RE-routers. However, at least a spanning tree has to be contained in any solution (since every node requests traffic from any other node).

Fig. 8– Fig. 13 show the trade-off between energy savings vs. the value of Γ for each pair of $(\bar{\gamma}, \hat{\gamma})$. The percentage of energy savings of the GreenRobustRE is computed in comparison with the case, that all links on the network are active (non-EAR solution). In both the Abilene and the Germany17 network, the solutions do not change when $\Gamma \geq \frac{|D|}{2}$, and thus the x-axis is cut at $\Gamma = \frac{|D|}{2}$. Clearly, a high value of Γ reduces the amount of energy savings for the network. From a technical point of view, an increase in Γ leads to higher compression multipliers in (3) which are directly linked to bigger coefficients in the same constraint. Thus, more capacity is needed and energy consumption increases (potential energy savings decrease).

However, we observe that the energy savings are only reduced at low values of Γ . The energy level becomes constant after a certain level of robustness is requested. For example in Fig. 8 the amount of energy savings does not change when $\Gamma \geq 10$, $\Gamma \geq 30$ and $\Gamma \geq 40$, respectively. Similar observations can be drawn from Fig. 9 – Fig. 13. An explanation of this phenomenon can be found in the distribution of the demand volumes. A fraction of the demands is dominating the others in volume. Hence, when the value of Γ covers all of these dominating demands, increasing Γ does not affect the routing solution and the percentage of energy savings remains stable. Fig. 10 ($\bar{\gamma} = 0.8, \hat{\gamma} = 0.05$) shows the extreme case, where the solution is already fully robust for $\Gamma = 0$, i.e., it is identical to the solution of $\Gamma = |D|$. This means the routing for a certain Γ has enough (spare) capacities to cover additional fluctuations without employing more links / RE routers.

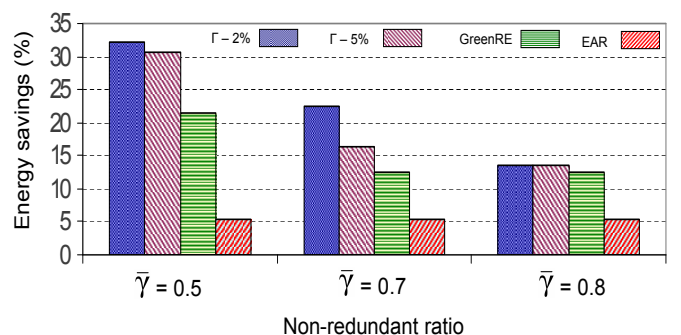


Fig. 6. Abilene network - GreenRobustRE vs. GreenRE vs. Classical EAR.

2) *GreenRobustRE vs. GreenRE vs. Classical EAR*: In this section, we exemplarily show that GreenRobustRE outperforms GreenRE and the classical EAR in case only a

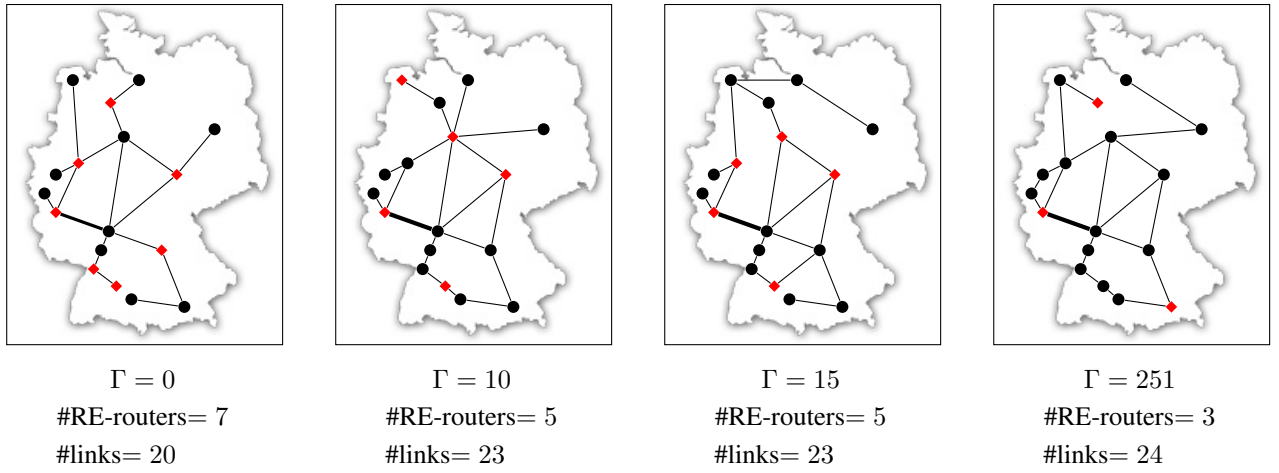


Fig. 4. Routing Solutions on Germany17, $\bar{\gamma} = 0.7, \hat{\gamma} = 0.2$

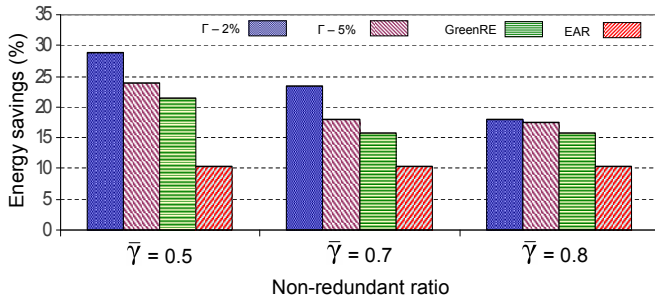


Fig. 7. Germany17 - GreenRobustRE vs. GreenRE vs. Classical EAR.

few traffic pairs deviate their redundancy elimination values simultaneously. Bertsimas and Sim [7] proved that already for small values of Γ , the capacity constraints are not violated with high probability. Experiments in [18] for uncertain demands confirm and even strengthen this result (but $\Gamma = 0$ yields infeasible solutions over almost realizations). Therefore, we compare EAR and GreenRE with Γ as low as 2% and 5% of the traffic pairs. Since we do not know the fluctuation of redundancy elimination in the (deterministic) GreenRE model, γ^{st} needs to be underestimated by the worst case realization, i.e., $\gamma^{st} = \bar{\gamma}^{st} + \hat{\gamma}^{st}$. Note that by this choice of γ^{st} , the GreenRE is equivalent to the GreenRobustRE model with $\Gamma = |D|$.

The estimated values of unique traffic and its deviation used in this section are $(\bar{\gamma} = 0.5, \hat{\gamma} = 0.25)$, $(\bar{\gamma} = 0.7, \hat{\gamma} = 0.2)$ and $(\bar{\gamma} = 0.8, \hat{\gamma} = 0.1)$. On the x-axis of Fig. 6 and Fig. 7, four columns for each value of $\bar{\gamma}$ represent the GreenRobustRE (with $\Gamma = 2\%$ and 5% of the total traffic pairs), the GreenRE, and the classical EAR.

We observe that the lowest energy savings are achieved by EAR. The energy efficiency for the network is improved when combining redundancy elimination and EAR (GreenRE). More importantly, the GreenRobustRE always outperforms GreenRE, in many cases even by a considerable amount. Referring to Section IV-B1, the GreenRobustRE model converges

to the GreenRE model if the robustness level is increased (e.g. more than 50% of the traffic pairs).

Altogether, we observe that in cases where a worst-case analysis is not necessary, but rather the congestion should be avoided with high probability, the Γ robustness approach yields a significant improvement over previously proposed models. By the GreenRobustRE model, network operators can draw more accurate estimations (both in quality and feasibility) of energy savings for their network depending on the level of desired robustness. In this context, solutions for different Γ can support a well-reasoned decision making.

V. CONCLUSION

In this paper we have proposed a concept for embedding data uncertainty into state-of-the-art models for minimizing energy consumption of backbone networks. Therefore, we formally define the GreenRobustRE problem and model it using robust optimization and MILP. Taking traffic redundancy uncertainties into account, the GreenRobustRE model provides an accurate model for potential energy savings in backbone networks. Based on a case study with real-life traffic demands, we show the relation between energy savings and the desired robustness for the network. Further, we give insights in the relation between this and earlier proposed models, showing that the GreenRobustRE model is clearly superior in our test-cases.

Future work can be splitted into three categories. At first, the strength of the GreenRobustRE model should be underlined by a bigger and a more detailed computational evaluation. Second (and related to the first point), more sophisticated solution approaches should be developed. While the standard MIP solution algorithm allows for a straight-forward optimal solutions in smaller examples, in bigger examples only primal and dual bounds can be derived. Therefore, we want to expand on either a fine-grained tuning of the MIP (i.e., by cutting planes) and/or propose a heuristic planning algorithm. Thus, we refer to [20], where similar work is done for the basic GreenRE model and which, we believe, could be extended

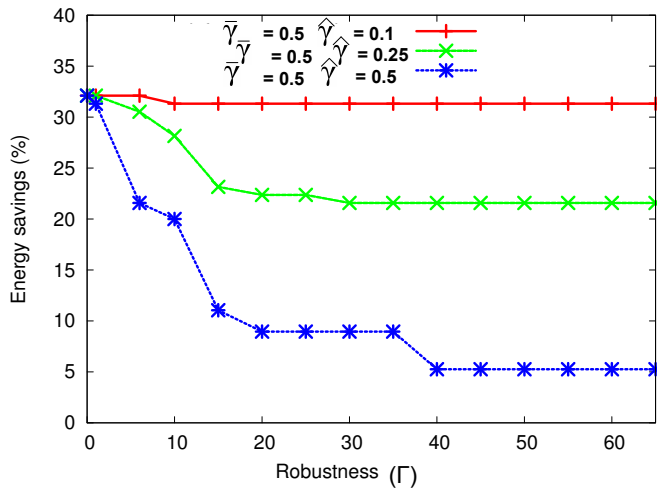


Fig. 8. Energy savings vs. Γ -robustness for Abilene network ($\bar{\gamma} = 0.5$).

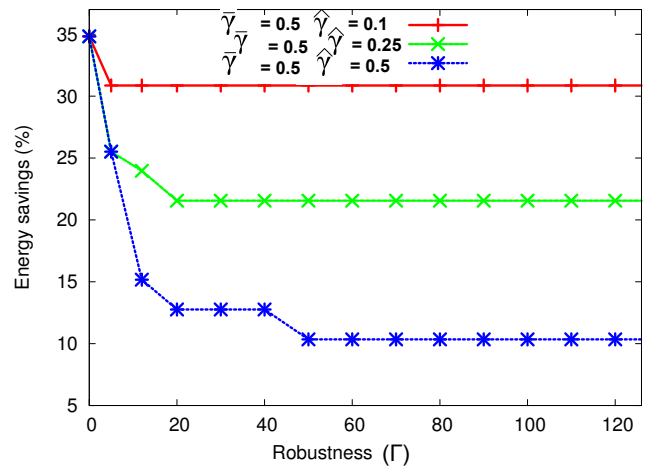


Fig. 11. Energy savings vs. Γ -robustness for Germany17 network ($\bar{\gamma} = 0.5$).

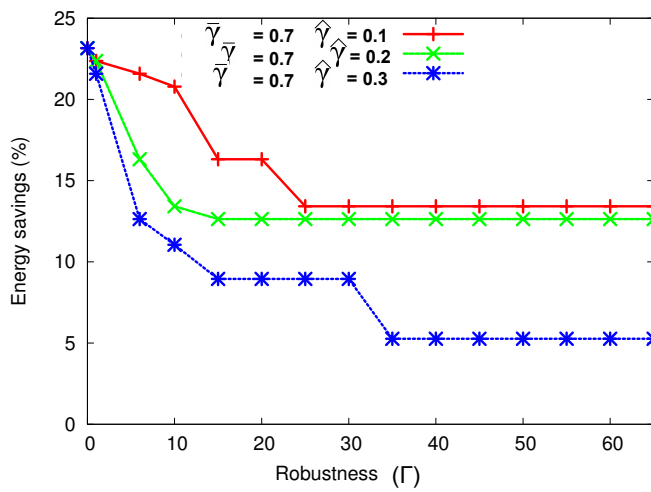


Fig. 9. Energy savings vs. Γ -robustness for Abilene network ($\bar{\gamma} = 0.7$).

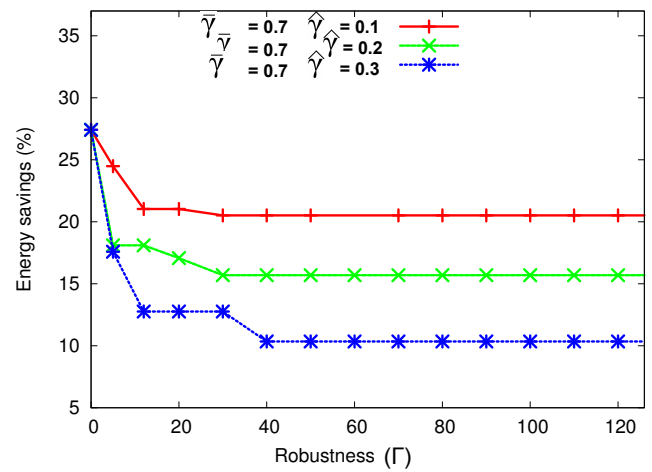


Fig. 12. Energy savings vs. Γ -robustness for Germany17 network ($\bar{\gamma} = 0.7$).

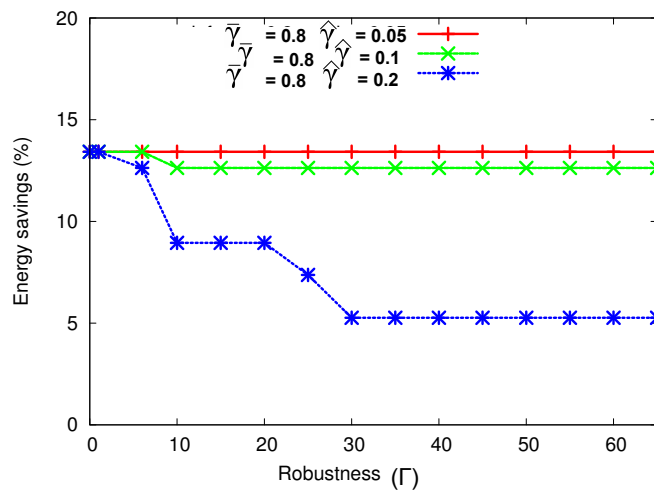


Fig. 10. Energy savings vs. Γ -robustness for Abilene network ($\bar{\gamma} = 0.8$).

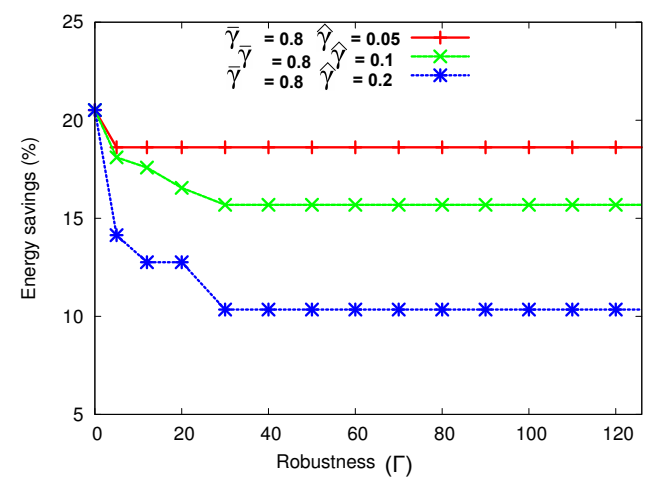


Fig. 13. Energy savings vs. Γ -robustness for Germany17 network ($\bar{\gamma} = 0.8$).

to the GreenRobustRE model as well (similar to [19] for uncertain demand values).

Thirdly, our model can be expanded to include more general data uncertainties, i.e., demand fluctuations should be considered as well. Clearly, for such a model, the two points above need to be applied as well.

ACKNOWLEDGMENT

This work has been done in collaboration with the Network Boost Project (Orange Labs), was partly funded by the ANR DIMAGREEN, and promoted by the PROCOPE program, a bilateral program funded by the German Academic Exchange Service (DAAD) and by the French Ministry of Foreign Affairs. The authors would like to thank Joanna Moulhierac from the COATI project for her advices and support.

REFERENCES

- [1] A. Altin, E. Amaldi, P. Belotti, and M. C. Pinar, "Provisioning Virtual Private Networks under Traffic Uncertainty," in *Networks*, vol. 49, 2007, pp. 100–115.
- [2] A. Anand, A. Gupta, A. Akella, S. Seshan, and S. Shenker, "Packet Caches on Routers: the Implications of Universal Redundant Traffic Elimination," in *Proc. of ACM SIGCOMM*, 2008, pp. 49 – 71.
- [3] A. Anand, C. Muthukrishnan, A. Akella, and R. Ramjee, "Redundancy in Network Traffic: Findings and Implications," in *Proc. of ACM SIGMETRICS*, 2009, pp. 37 – 48.
- [4] A. Anand, V. Sekar, and A. Akella, "SmartRE: an Architecture for Coordinated Network-wide Redundancy Elimination," in *Proc. of ACM SIGCOMM*, 2009, pp. 87 – 98.
- [5] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*, ser. Princeton Series in Applied Mathematics, 2009.
- [6] A. Ben-Tal and A. Nemirovski, "Robust Optimization - Methodology and Application," in *Math. Prog.*, vol. 92, 2002, pp. 453 – 480.
- [7] D. Bertsimas and M. Sim, "Robust Discrete Optimization and Network Flows," in *Math. Prog.*, vol. 98, 2003, pp. 49 – 71.
- [8] —, "The Price of Robustness," in *Operations Research*, vol. 52, 2004, pp. 35 – 53.
- [9] "BlueCoat: WAN Optimization," <http://www.bluecoat.com/>.
- [10] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsang, and S. Wright, "Power Awareness in Network Design and Routing," in *Proc. of IEEE INFOCOM*, 2008, pp. 457 – 465.
- [11] L. Chiaraviglio, M. Mellia, and F. Neri, "Minimizing ISP Network Energy Cost: Formulation and Solutions," in *IEEE/ACM Trans. Netw.*, vol. 20, 2011, pp. 463 –476.
- [12] S. Duhovniko, A. M. C. A. Koster, M. Kutschka, F. Rambach, and D. Schupke, "Γ-Robust Network Design for Mixed-Line-Rate-Planning of Optical Networks," in *Proc. Optical Fiber Communication - National Fiber Optic Engineers Conference (OFC/NFOEC)*, 2013.
- [13] F. Giroire, D. Mazauric, J. Moulhierac, and B. Onfroy, "Minimizing Routing Energy Consumption: from Theoretical to Practical Results," in *Proc. of IEEE/ACM GreenCom*, 2010, pp. 252 – 259.
- [14] F. Giroire, J. Moulhierac, T. K. Phan, and F. Roudaut, "Minimization of Network Power Consumption with Redundancy Elimination," in *Proc. of IFIP NETWORKING*, 2012, pp. 247 – 258.
- [15] T. J. Grevers and J. Christner, "Application Acceleration and WAN Optimization Fundamentals," in *Cisco Press*, 2007.
- [16] M. Gupta and S. Singh, "Greening of The Internet," in *Proc. of ACM SIGCOMM*, 2003, pp. 19 – 26.
- [17] IBM – ILOG, "IBM ILOG CPLEX Optimizer, version 12.4," 2012, <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>.
- [18] A. M. C. A. Koster, M. Kutschka, and C. Raack, "On the Robustness of Optimal Network Designs," in *Proc. of IEEE ICC*, 2011, pp. 1 – 5.
- [19] A. M. C. A. Koster, M. Kutschka, and C. Raack, "Robust Network Design: Formulations, Valid Inequalities, and Computations," *Networks*, vol. 61, no. 2, pp. 128–149, 2013.
- [20] A. M. C. A. Koster, T. K. Phan, and M. Tieves, "Extended cutset inequalities for the network power consumption problem," *Electronic Notes in Discrete Mathematics*, vol. 41, no. 0, pp. 69 – 76, 2013.
- [21] P. Mahadevan, P. Sharma, and S. Banerjee, "A Power Benchmarking Framework for Network Devices," in *Proc. of IFIP NETWORKING*, 2009, pp. 795 – 808.
- [22] S. Orłowski, R. Wessälly, M. Pióro, and A. Tomaszewski, "SNDlib 1.0–Survivable Network Design Library," *Networks*, vol. 55, pp. 276–286, 2010, <http://sndlib.zib.de>.
- [23] "Riverbed: WAN Optimization," http://www.riverbed.com/us/solutions/wan_optimization/.
- [24] Y. Song, K. Guo, and L. Gao, "Redundancy-Aware Routing with Limited Resources," in *Proc. of IEEE ICCCN*, 2010, pp. 1 – 6.
- [25] N. T. Spring and D. Wetherall, "A Protocol-Independent Technique for Eliminating Redundant Network Traffic," in *Proc. of ACM SIGCOMM*, 2000, pp. 87 – 95.
- [26] M. Zhang, C. Yi, B. Liu, and B. Zhang, "GreenTE: Power-aware Traffic Engineering," in *Proc. of IEEE ICNP*, 2010, pp. 21 – 30.