GreenCoop: Cooperative Green Routing with Energy-efficient Servers

Luca Chiaraviglio
Dipartimento di Elettronica
Politecnico di Torino - Italy
chiaraviglio@tlc.polito.it

Ibrahim Matta
Computer Science Department
Boston University - USA
matta@cs.bu.edu

ABSTRACT

Energy-efficient communication has recently become a key challenge for both researchers and industries. In this paper, we propose a new model in which a Content Provider and an Internet Service Provider *cooperate* to reduce the total power consumption. We solve the problem optimally and compare it with a classic formulation, whose aim is to minimize user delay. Results, although preliminary, show that power savings can be huge: up to 71% on real ISP topologies. We also show how the degree of cooperation impacts overall power consumption. Finally, we consider the impact of the Content Provider location on the total power savings.

1. INTRODUCTION

Energy-efficient communication has become a sensible problem in the last few years. According to recent studies [1] the next century will be dominated by environmental changes due to global warming, and key players in this process will be the production, distribution and the use of energy. Therefore, actions to reduce power consumption and improve energy efficiency are becoming imperative.

Considering the Information and Communication Technology sector (ICT), current estimates [2] show that ICT is responsible for 1% to 10% of the worldwide energy consumption, and this trend is expected to grow even more in the future. Starting from the seminal work of [3], several approaches have been proposed in order to reduce power consumption of ICT. For example, in [4] the authors consider the minimization of the power consumed by an Internet Service Provider (ISP) network. Moreover, new recent approaches like [5] aim at reducing the power consumption of big Content Providers (CP) considering the variation in electricity prices. Nevertheless, none of these previous studies considers the minimization of the *total* power consumption of both ISP and CP, so that great power savings can be achieved by considering a mutual objective.

In [6] the authors solve jointly the traffic engineering and content distribution problem, showing that great improvements in Quality of Service (QoS) can be obtained if CP and ISP pursue the same objective. However, this solution does not consider power consumption at all and consequently the power waste can be huge.

In this work we consider a design problem in which both CP and ISP *cooperate* in order to reduce overall power consumption. In particular, we model the system as an optimization problem in which the objective function is the minimization of the total power consumed by the CP and the ISP, subject to a user delay constraint. We test our model on real ISP topologies and considering realistic power figures. Our preliminary results show that there is ample room to pursue a cooperative green approach, so that large power savings can be easily achieved.

While the proposed approach shows that there is a great opportunity to save power through cooperation, it is also true that CP and ISP are not willing to share sensible data such as the network topology or end-to-end traffic demand. We are therefore encouraged to pursue further research in this direction in order to improve our model by limiting the amount of shared information among CP and ISP.

2. PROBLEM FORMULATION

The main goal of our approach is to minimize the power consumption jointly between the CP and the ISP. In particular, we assume that the ISP is the owner of the network infrastructure, so that it manages a physical topology, i.e. a set of nodes and links. The CP instead is composed of a number of servers connected to the ISP. When a user (terminal) asks for a CP's resource, we assume that the resource is replicated over the CP infrastructure, so that the user can be potentially served by *any* of the servers of the CP.

More formally, we represent the ISP topology as a digraph $G = (V, E)$, where V is the set of vertices and E is the set of edges. Vertices represent network nodes, while edges represent network links. We denote by $N = |V|$ and $L = |E|$ the total number of nodes and links, respectively. Let C_l be the capacity of link l, and let $U_l^{MAX} \in [0, 1]$ be the maximum link utilization that can be tolerated [1]. S is the set of servers of the content provider. Denote by W_s the maximum load allowed on server $s \in S$. Let R_t be the traffic demand between terminal $t \in T$ and the content provider S. Moreover, let x^{st} be continuous variables representing the amount of traffic between a source node s and a terminal t. We divide x^{st} into x_{cp}^{st} and x_{bg}^{st} to denote the amount of traffic originating from the content provider under consideration

---

[1] Link utilization is normally kept below 100% due to QoS requirements.

and from other content providers, respectively. Actually, x_{bg}^{st} are constants, i.e. the considered CP can not modify these traffic demands. On the other hand, we assume that x_{cp}^{st} are continuous variables so that a traffic demand R_t from terminal t can be served by anyone of the $s \in S$ CP servers, while satisfying load and delay constraints.

In order to describe the network topology, we can either choose between a *node-link* formulation or a *path-link* formulation [7]. In particular, in the node-link formulation each node computes the flow conservation law, while in the path-link formulation the flow conservation is done on the entire path. We decided to choose the path-link formulation since: i) it requires a lower number of variables than the node-link formulation ($O(N^2)$ instead of $O(N^3)$); ii) ISP can easily control the paths since they are pre-computed prior to launching the problem; iii) additional constraints, such as a minimum number of disjoint paths for each source and destination pair, can be easily enforced.

More formally, let δ_{lp}^{st} be constants which take the values of 1 if link l belongs to path p carrying demand from s to t, 0 otherwise. Let z_p^{st} and q_p^{st} be continuous variables representing the amount of traffic from s to t on path p for the considered CP and for other CPs, respectively. Let $\mathcal{P}(s,t)$ the set of pre-computed paths from s to t. Additionally, let f_l be the total amount of flow on link l, split into f_l^{cp} and f_l^{bg} for the considered CP and for other CPs. Let D_l be the delay on link l, which can be modeled as $D_l = 1/(C_l - f_l) + D_l^p$, where $1/(C_l - f_l)$ and D_l^p denote the queueing delay and the propagation delay of link l, respectively.

Finally, we consider the power consumption of devices. Let y_n and y_s be binary variables which take the value of 1 if node n and server s are powered on, respectively. $P_n^c y_n$ and $P_s^c y_s$ represent the static amount of power consumed by node n and server s when powered on. Moreover, let $P_l^d(f_l)$, $P_n^d(\sum_{l \in \mathcal{L}(n)} f_l)$, $P_s^d(\sum_{t \in T} x_{cp}^{st})$ be monotonically increasing convex functions representing the dynamic power consumption of link l, node n and server s, respectively. $\mathcal{L}(n)$ denotes the set of links incident to node n.

Given the previous notations, we recall the classic formulation of [6] and then present a new green formulation.

2.1 Classical Design

The objective of the classic design problem presented in [6] is the minimization of the overall delay experienced by all users in the network. The problem is formalized as follows:

ClassicCoop $\quad \min \sum_l f_l^{cp} D_l \quad$ subject to:

$$\sum_{s \in S} x_{cp}^{st} = R_t \quad \forall t \in T \tag{1}$$

$$\sum_{t \in T} x_{cp}^{st} \leq W_s \quad \forall s \in S \tag{2}$$

$$\sum_{p \in \mathcal{P}(s,t)} q_p^{st} = x_{bg}^{st} \quad \forall s,t \tag{3}$$

$$x_{cp}^{st} = \sum_{p \in \mathcal{P}(s,t)} z_p^{st} \quad \forall s,t \tag{4}$$

$$f_l^{cp} = \sum_{s,t,p \in \mathcal{P}(s,t)} \delta_{lp}^{st} z_p^{st} \quad \forall l \in E \tag{5}$$

$$f_l^{bg} = \sum_{s,t,p \in \mathcal{P}(s,t)} \delta_{lp}^{st} q_p^{st} \quad \forall l \in E \tag{6}$$

$$f_l = f_l^{bg} + f_l^{cp} \leq C_l U_l^{MAX} \quad \forall l \in E \tag{7}$$

Control variables: $z_p^{st} \geq 0$, $q_p^{st} \geq 0$.

More precisely, Eq.1 ensures that the total traffic demand R_t from terminal t is served by the CP, Eq.2 limits the total traffic on each server by the maximum load, Eq.3 splits the traffic demand x_{bg}^{st} among the different paths, Eq.4 computes the total amount of traffic from each server to each terminal of the CP. Finally, Eq.5-7 compute the total flow on each link and impose the maximum link utilization constraint. This formulation does not take into account power consumption at all: we therefore compute the power consumption of devices as a post-processing phase.

ClassicCoop falls into the class of convex optimization problems, for which finding a local optimum is equivalent to finding the global optimum.

2.2 A New Green Cooperation

In order to consider the power consumption we propose a novel approach in which CP and ISP share information to minimize the global power consumption. In particular, the problem can be formalized as follows:

GreenCoop $\quad \min(P_{TOT} = P_{CP} + P_{ISP}) \quad$ s.t.:

$$P_{CP} = \sum_{s \in S} P_s^d \left(\sum_{t \in T} x_{cp}^{st} \right) + P_s^c y_s \tag{8}$$

$$P_{ISP} = \sum_{l \in E} P_l^d(f_l) + \sum_{n \in V} \left[P_n^d \left(\sum_{l \in \mathcal{L}(n)} f_l \right) + P_n^c y_n \right] \tag{9}$$

$$\sum_{s \in S} x_{cp}^{st} = R_t \quad \forall t \in T \tag{10}$$

$$\sum_{t \in T} x_{cp}^{st} \leq W_s \quad \forall s \in S \tag{11}$$

$$\sum_{p \in \mathcal{P}(s,t)} q_p^{st} = x_{bg}^{st} \quad \forall s,t \tag{12}$$

$$x_{cp}^{st} = \sum_{p \in \mathcal{P}(s,t)} z_p^{st} \quad \forall s,t \tag{13}$$

$$f_l^{cp} = \sum_{s,t,p \in \mathcal{P}(s,t)} \delta_{lp}^{st} z_p^{st} \quad \forall l \in E \tag{14}$$

$$f_l^{bg} = \sum_{s,t,p \in \mathcal{P}(s,t)} \delta_{lp}^{st} q_p^{st} \quad \forall l \in E \tag{15}$$

$$f_l = f_l^{bg} + f_l^{cp} \leq C_l U_l^{MAX} \quad \forall l \in E \tag{16}$$

$$d_l \geq a_i f_l + b_i \quad \forall l, i \in I \tag{17}$$

$$\frac{\sum_l d_l}{|T|} \leq D^{MAX} \tag{18}$$

$$\sum_{t \in T} x_{cp}^{st} \leq M_s y_s \quad \forall s \in S \tag{19}$$

$$\sum_{l \in \mathcal{L}(n)} f_l \leq M_n y_n \quad \forall n \in V \tag{20}$$

Control variables: $z_p^{st} \geq 0$, $q_p^{st} \geq 0$, $y_n \in \{0,1\}$, $y_s \in \{0,1\}$.

Eq.8-9 compute the total power for the CP and the ISP, respectively. Eq.10 guarantees the traffic demand constraint. Eq.11 limits the maximum load on each server. Routing constraints are specified by Eq.12-13. The total flow on each link is computed and constrained by Eq.14- 16. Then, Eq.17 computes the total delay for each link, using the additional variables $d_l \geq 0$. Here the delay function $D_l = 1/(C_l - f_l) + D_l^p$ is approximated by I linear segments. Eq.18

Table 1: Power Consumption Model [W]

P_n^c	$P_l^d(f_l)$	$P_n^d\left(\sum_{l\in\mathcal{L}(n)}f_l\right)$	P_s^c	$P_s^d\left(\sum_{t\in T}x_{cp}^{st}\right)$
100	$20f_lA_l$	$20\sum_{l\in\mathcal{L}(n)}f_l$	200 ± 100	$(40\pm20)\sum_{t\in T}x_{cp}^{st}$

bounds the average delay of users. Finally, Eq.19-20 impose powering-on a network node and a server, respectively, if their incoming/outgoing flows are larger than zero, adopting a big-M method, i.e. $M_s \geq W_s$ and $M_n \geq \sum_{l\in\mathcal{L}(n)}C_l$.

GreenCoop falls into the class of mixed-integer linear problems, for which optimal solutions can be found for example through the branch-and-bound algorithm.

3. PERFORMANCE EVALUATION

We test the effectiveness of the model proposed using ISP backbone topologies obtained from RocketFuel [8]. The topologies are first pre-processed using a simple shortest path algorithm to obtain the set of paths, using the measured weights as link costs. In particular, for each (s,t) we compute up to two completely disjoint paths. This reflects normal behaviour of ISP which guarantees alternate paths for failure protection. Moreover, the capacity C_l is set to 10 Gbps for each link, since the topologies considered are those of tier-1 ISPs. Links are utilized up to 50% of their capacity, i.e. $U_l^{MAX} = 0.5 \quad \forall l \in E$, to avoid congestion and to guarantee QoS. We assume that nodes are connected by optical links, in which the optical carrier is regenerated by amplifiers. For each link we randomly assign a number of amplifiers A_l uniformly distributed between 1 and 5.

Considering the CP, the maximum load on each server is set to the network capacity offered at that node, i.e. $W_s = \sum_{l\in\mathcal{L}(s)}C_lU_l^{MAX}$. We consider the case in which the CP infrastructure is composed of 15 servers, adopting different strategies for server placement over the ISP topology. The CP traffic demand R_t is modeled according to a Pareto distribution, with a variable lower bound R_t^{min} and a constant upper bound R_t^{MAX} given by the total capacity offered at that node, i.e. $R_t^{MAX} = \sum_{l\in\mathcal{L}(t)}C_lU_l^{MAX}$. Unless otherwise specified, $D^{MAX} = 300$ ms and $R_t^{min} = 100$ Mbps.

Tab.1 describes the model used to evaluate the power consumption. Here we are assuming next-generation devices able to adapt their power with traffic flow. Considering the ISP, the power consumption of nodes is composed of a constant term P_n^c due to the chassis static power plus an additional term P_n^d which scales linearly with traffic flow. The constant values are extracted by interpolating the power measurements of real devices under high load [9]. Moreover, the power consumption of a link P_l^d depends linearly on both the load and the number of amplifiers A_l between nodes, as reported in [10].

Focusing on CP, the server power consumption is also modeled by a static term P_s^c and a dynamic term P_s^d: in this case instead the slope is higher due to the presence of backup elements and power supplies, which actually double the server power consumption. Moreover, an additional random variation of 50% in the server power is introduced to model energy price fluctuation as reported in [5]. For the sake of simplicity we do not consider any additional background traffic of other CPs, since our goal is mainly to assess the maximum power savings achievable by the whole system composed of the ISP and the considered CP. Finally, 50% of randomly chosen nodes are selected as terminals t.

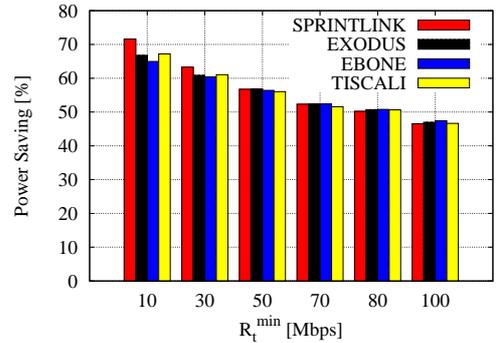In the following sections we investigate potential power

Figure 1: Average power saving against traffic variation for different topologies.

savings under different strategies for server placement over the topologies, considering also the impact of variation in the constraints R_t^{min} and D^{MAX}.

3.1 Preferential Server Assignment

In this set of experiments we assign the servers of the considered CP to nodes using a preferential degree placement. ISP nodes are first grouped according to the city in which they are located, and then the groups are sorted by decreasing number of links with other cities. Finally, the CP servers are assigned to the cities with the highest connection degree, one CP server per city. A random node inside each selected city is chosen as server location.

We solve optimally both ClassicCoop and GreenCoop, considering different scenarios. Fig.1 details the power savings against increasing R_t^{min}, for different topologies. Values are averaged over 10 runs, where for each run we chose a different set of terminal nodes. Astonishingly, GreenCoop is able to save from 47% to more than 65% of power relative to ClassicCoop, for all the topologies considered, with a maximum power saving of 71% for the SprintLink topology. This is due to the fact that the green model explicitly takes into account the power consumption of devices, powering-on only the minimal amount of resources needed to satisfy the traffic demand. Moreover, the savings are decreasing as the traffic increases, suggesting that a larger number devices need to be powered on to satisfy the constraints.

To give more insight, we add a parameter α to weigh differently the CP and the ISP power, so that the objective function of GreenCoop becomes $\alpha P_{CP} + (1-\alpha)P_{ISP}$ and $\alpha \in [0,1]$. With $\alpha = 0$ just the minimization of the power within the ISP infrastructure is pursued, while with $\alpha = 1$ only the optimization of the content provider is considered. Intuitively, a CP-only power optimization may result in choosing a more energy-efficient server but so much more distant that more power is consumed over the network. Conversely, an ISP-only power optimization may yield a server that is closer to the terminal but whose power cost is high. Thus, a joint power optimization should provide the right balance and yield higher power savings.

Fig.2 details the total power consumption as α varies, considering the different topologies. In this case the best savings can be obtained only when both the CP and ISP power are taken into account. For example, the total power of the SprintLink topology is more than 41kW with $\alpha = 1$, decreasing to 23kW with $\alpha = 0.5$. In this case the largest part of power consumption is due to ISP, so that the power due
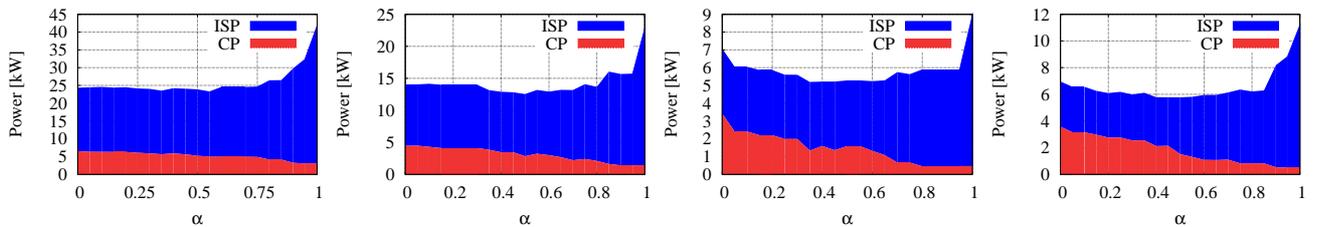
Figure 2: Power consumption components versus α. From left to right: SPRINTLINK, EXODUS, EBONE and TISCALI topologies.
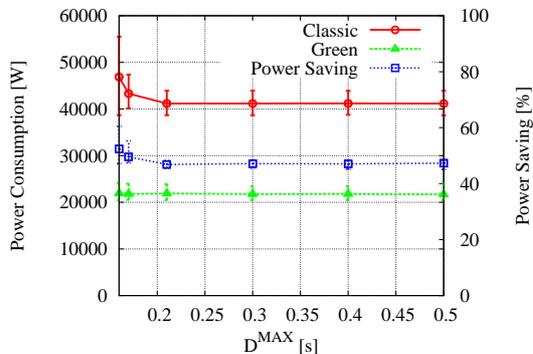

Figure 3: Maximum admissible delay for the SPRINTLINK topology.

to CP does not increase significantly as α decreases further. Instead, if we consider for example the EBone topology, the total power is clearly minimized only when $\alpha = 0.5$. The intuition suggests that a great amount of power is wasted if the ISP and CP individually optimize their own power consumption, while less power is required if both of them jointly pursue the minimization of the total power of the entire system, i.e. $\alpha \approx 0.5$.

In order to assess the impact of QoS constraints, we consider the variation in the maximum delay D^{MAX}. Fig.3 reports the power consumption for both ClassicCoop and GreenCoop for different D^{MAX} values, considering the Sprint-Link topology[2]. Values are averaged over 10 different runs, while error bars report minimum and maximum values. Interestingly, with ClassicCoop the power grows notably with decreasing D^{MAX}, while GreenCoop consumes around the same amount of power even for small D^{MAX}. The intuition suggests that GreenCoop can more wisely adapt to stricter QoS constraints, leading to higher power savings (around 52%) even for small D^{MAX}.

In the last part of our work we adopt a completely random node selection for placing the CP servers. Due to lack of space we refer the reader to [11] for these results. In brief, a random placement can waste a considerable amount of power.

4. CONCLUSIONS AND FUTURE WORK

In this paper we have proposed GreenCoop, a new model in which both CP and ISP cooperate to reduce the overall power consumption. Our preliminary results on real topologies show that large power savings are achievable, up to 71% relative to a classic formulation of the joint problem.

---

[2]Here we have enforced the maximum admissible delay constraint also for the ClassicCoop model.

We recognize that this work is a first step towards a comprehensive approach. As future work, we first plan to develop an energy-efficient distributed algorithm to reduce the amount of shared information between CP and ISP. Another possible extension is the interaction of multiple CPs over the ISP to minimize power consumption, considering also the effects of server virtualization on the potential power savings possible with many CPs. Finally, different classes of QoS can be taken into account.

5. ACKNOWLEDGMENTS

6. REFERENCES

[1] A.P. Sokolov et al., "Probabilistic Forecast for 21st Century Climate Based on Uncertainties in Emissions (without Policy) and Climate Parameters", *MIT Global Change Program Report*, Cambridge, USA, January 2009.

[2] SMART 2020, *Enabling the low carbon economy in the information age*, http://www.theclimategroup.org, 2009.

[3] M. Gupta, S. Singh, "Greening of the Internet", *ACM SIGCOMM*, Karlsruhe, Germany, August 2003.

[4] P. Barford, J. Chabarek, C. Estan, J. Sommers, D. Tsiang and S. Wright, "Power Awareness in Network Design and Routing," *IEEE INFOCOM*, Phoenix, USA, April 2008.

[5] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag and B. Maggs., "Cutting the Electric Bill for Internet-Scale Systems," *ACM SIGCOMM*, Barcelona, Spain, 2009.

[6] W. Jiang, R. Zhang-Shen, J. Rexford, and M. Chiang, "Cooperative Content Distribution and Traffic Engineering in an ISP Network," *ACM SIGMETRICS*, Seattle, Washington, USA, June 2009. vspace1.5mm

[7] M. Pioro, D. Medhi, *Routing, Flow and Capacity Design in Communication and Computer Networks*, The Morgan Kaufmann Series in Networking, New York, USA, 2004.

[8] N. Spring, R. Mahajan and D. Wetherall, "Measuring ISP Topologies with Rocketfuel," *ACM SIGCOMM*, Pittsburgh, PA, August 2002.

[9] E. Bonetto, L. Chiaraviglio, D. Cuda, G. Gavilanes and F. Neri, "Optical Technologies Can Improve the Energy Efficiency of Networks," *ECOC 2009*, Vienna, Austria, September 2009.

[10] L. Chiaraviglio, M. Mellia, F. Neri, "Energy-aware Backbone Networks: a Case Study," *IEEE GreenComm'09*, Dresden, Germany, June 2009.

[11] L. Chiaraviglio, I. Matta, "GreenCoop: Cooperative Green Routing with Energy-efficient Servers," *Technical Report BUCS-TR-2009-034*, CS Department, Boston University, November 2009.