

Energy-Aware Routing in Multiple Domains Software-Defined Networks



Adriana Fernández-Fernández^a, Cristina Cervelló-Pastor^b,
and Leonardo Ochoa-Aday^c

^{a,b,c}Department of Network Engineering, Universitat Politècnica de Catalunya (UPC), Esteve Terradas 7, Castelldefels, 08860, Spain

^{a,c}Student Member, IEEE ^bMember, IEEE

adriana.fernandez@entel.upc.edu, cristina@entel.upc.edu, leonardo.ochoa@entel.upc.edu

KEYWORD

ABSTRACT

Distributed Routing Algorithm, Software-Defined Networks, In-band Control Traffic, Energy-Aware Routing, Traffic Engineering

The growing energy consumption of communication networks has attracted the attention of the networking researchers in the last decade. In this context, the new architecture of Software-Defined Networks (SDN) allows a flexible programmability, suitable for the power-consumption optimization problem. In this paper we address the issue of designing a novel distributed routing algorithm that optimizes the power consumption in large scale SDN with multiple domains. The solution proposed, called DEAR (Distributed Energy-Aware Routing), tackles the problem of minimizing the number of links that can be used to satisfy a given data traffic demand under performance constraints such as control traffic delay and link utilization. To this end, we present a complete formulation of the optimization problem that considers routing requirements for control and data plane communications. Simulation results confirm that the proposed solution enables the achievement of significant energy savings.

1. Introduction

The high energy consumption generated by network elements and the expansion of Internet, have brought power consumption of data networks to the forefront as a major optimization concern (Van Heddeghem et al., 2014). According to (Gelenbe and Caseau, 2015) in 2012, close to 4.7% of the world's electrical energy was consumed by ICT, releasing into the atmosphere roughly 1.7% of the total CO₂ emissions. Moreover, recent studies state that energy demand of ICT sector is growing faster than the overall one (Aebischer and Hilty, 2015).

Among the main ICT sectors, telecommunication networks account for more than a third part of the ICT total energy consumption (The Climate Group, 2008). As a result, the reduction of power consumption in Internet Service Provider (ISP) backbone networks is crucial to accomplish significant energy savings in this sector. For this problem, the emerging paradigm of Software-Defined Networks (SDN) can be seen as an attractive solution.

In SDN (Kreutz et al., 2015), control functions are decoupled from forwarding devices and are logically centralized in a new entity called controller. The controller has a global network view and can manage network tasks without the need of additional software in each of the switching elements. Meanwhile, the network devices only forward traffic according to the rules set by the controller. In this paper we address the problem of optimizing the power consumption in OpenFlow networks (McKeown et al., 2008).

The idea of saving energy by turning off unused networks elements such as line cards or port interfaces, was first considered by (Gupta and Singh, 2003). In (Zhang et al., 2010), the exact optimization problem of maximizing the total power saving under Maximum Link Utilization (MLU) and network delay constraints in



traditional networks, is formulated.

The use of OpenFlow for this purpose has already been included in other research papers. The authors of (Wang et al., 2016) formulated an optimization problem for finding minimum-power network subsets in hybrid SDN. Giroire et al. (Giroire et al., 2014) proposed an energy-aware routing approach, taking into account the limited rule space of TCAM (Ternary Content Addressable Memory) in SDN devices. The authors of (Wang et al., 2014) provided two greedy algorithms for minimizing the power of integrated chassis and line-cards used. For this they considered an expanded network topology according to the connections between the forwarding devices. However, all these related works considered a centralized approach.

In practice, the logically centralized control in SDN could be implemented with multiple distributed physical controllers, which is the scenario considered in this work. The hypothesis of our research is that in these scenarios, an effective optimization of power consumption could be achieved with a distributed energy-aware routing algorithm. Different from previous works, we focus on optimizing energy consumption in multiple domains OpenFlow networks with in-band control traffic.

The rest of this paper is structured as follows. In Section 2 we explain the main considerations of our distributed approach together with the network model considered and the mathematical formulation of our optimization model. The simulation strategies and the obtained results are discussed in Section 3. Finally, in Section 4 we conclude our work and outline future research guidelines.

2. Distributed Energy-Aware Routing Algorithm

The Distributed Energy-Aware Routing (DEAR) approach consists in the use of traffic engineering in each domain to optimize the overall power consumption. The idea is to find the routes between network elements that minimize the number of active links used to satisfy a given data data traffic demand, subject to the capacity constraint.

In order to ensure compatibility with SDN using in-band control traffic (Sharma et al., 2013), in this proposal control paths between controllers and switches are also established. This means that control messages are exchanged using the same links that data traffic without the need of additional links. This is a more realistic scenario for large backbone networks, where dedicated links to transfer the control messages between controllers and forwarding devices are impractical and cost-inefficient. In addition, to avoid additional traffic load in the controllers, we establish that data plane communications cannot be routed through these devices.

We consider a multiple domains SDN architecture, where each domain has a centralized controller with a number of predefined switches associated to it. We assume that each controller has a total knowledge of its domain topology and a partial knowledge of the global network topology, i.e., it has identified border nodes that it shares with each other domain. Inter-domain data traffic demands are routed in each domain using these nodes.

2.1 Network Model

Each controller domain is represented by a directed graph $G = (V, E)$, where V and E denote the set of nodes and links, respectively. Each link $e \in E$ has associated its capacity, denoted by c_e . The set $B = \{b_1, \dots, b_{|B|}\} \subset V$ contains the border nodes. D_v and D_w denote the set of intra-domain traffic flows for the data and control plane communications, respectively. D_u denote the set of inter-domain data traffic demands.

For each $k \in D_v$, let t_k denote its throughput and P_k be the set of paths that can be used to route this traffic. $P_c^k \subset P_k$ denote the set of paths that pass through the controller for each $k \in D_v$. Let $P_e^k \subset P_k$ be the set of paths that use link $e \in E$ for each $k \in D_v$. Similarly, it holds for D_w and D_u traffic flows.



2.2 Optimization Problem Formulation

The distributed proposal of our approach in multiple domains SDN, can be formulated as an Integer Linear Programming (ILP) model with two steps of optimization, using the following binary variables:

x_e : describes the state of a link $e \in E$.

$$x_e = \begin{cases} 1 & \text{if link } e \text{ is active,} \\ 0 & \text{otherwise.} \end{cases}$$

q_b^k : describes the selection of a border node b to route a traffic $k \in D_u$.

$$q_b^k = \begin{cases} 1 & \text{if border node } b \text{ is selected to route inter-domain traffic } k, \\ 0 & \text{otherwise.} \end{cases}$$

$l_{b,p}^k$: describes the selection of a path $p \in P_k$ to route a traffic $k \in D_u$ through border node b .

$$l_{b,p}^k = \begin{cases} 1 & \text{if path } p \text{ is selected to route inter-domain traffic } k \text{ through border node } b, \\ 0 & \text{otherwise.} \end{cases}$$

r_p^k : describes the selection of a path $p \in P_k$ to route a traffic $k \in D_v \cup D_w$.

$$r_p^k = \begin{cases} 1 & \text{if path } p \text{ is selected to route intra-domain traffic } k, \\ 0 & \text{otherwise.} \end{cases}$$

In the first step, each controller-instantiated agent individually computes the routing paths in its domain that minimize the number of links used. In this phase, performance constraints (e.g., control traffic delay and link utilization) could be included. Considering the notation of binary variables shown above, the optimization model of the first phase can be formulated as:

$$\text{minimize } \sum_{e \in E} x_e \quad (1)$$

subject to the following constraints:

$$\sum_{b \in B} q_b^k = 1 \quad \forall k \in D_u \quad (2)$$

$$\sum_{p \in P_k} l_{b,p}^k = q_b^k \quad \forall k \in D_u, \forall b \in B \quad (3)$$

$$\sum_{p \in P_k} r_p^k = 1 \quad \forall k \in D_v \cup D_w \quad (4)$$

$$r_p^k = 0 \quad \forall k \in D_v, \forall p \in P_c^k \quad (5)$$

$$l_{b,p}^k = 0 \quad \forall k \in D_u, \forall b \in B, \forall p \in P_c^k \quad (6)$$

$$\sum_{k \in D_u} \sum_{p \in P_c^k} \sum_{b \in B} l_{b,p}^k t_k + \sum_{k \in D_v \cup D_w} \sum_{p \in P_c^k} r_p^k t_k \leq c_e x_e \quad \forall e \in E \quad (7)$$

The objective function (1) minimizes the number of active links. Equation (2) ensures that exactly one border node is selected for every inter-domain data traffic demand. Equation (3) ensures that exactly one path is used to route every inter-domain data traffic demand through the border node selected. Equations (4) ensure that exactly one path is used to route every intra-domain traffic flow for the data and control plane communications. Equations (5) and (6) ensure that paths passing through the controller can not be used to route data plane communications. Equation (7) ensures that the total traffic in each active link $e \in E$ is less than its capacity c_e .

After completing this computation, the distributed control plane agents in different SDN domains must exchange some performance metric (e.g. MLU in each domain) and the identifier of the selected border nodes to route each inter-domain data traffic demand (i.e. $q_b^k \quad \forall k \in D_u$). The first element of this shared information is intended to be used as comparison metric to define the domain with the best performance, which is also the one with the lowest probability to run out of capacity, while the second one allows a proper and coherent rerouting of inter-domain data traffic demands.

In the second step, the agent of the domain with the best performance (less MLU, for instance) recomputes its energy-aware routing paths using now, for each inter-domain data traffic demand, the border nodes preselected by its neighbor domains. The corresponding problem for the second step of optimization could be formulated using these received identifiers in (3) of the model above.

3. Preliminary Results

In this section we describe the evaluation of our distributed approach and analyze the results obtained. We used the linear programming solver Gurobi Optimizer (Gurobi Optimization) to assess the performance of the ILP model. All computations were carried out on a computer equipped with 3.30 GHz Intel Core i7 and 16 GB RAM.

We conducted our simulations using a real network topology, Abilene, and the subset of online available traffic matrices measured on September 5th 2004 (Zhang, Y., 2004). The energy savings were computed as the number of links in sleep mode over the total amount of network links.




Figure 1: Percentage of shutdown links in the Abilene topology with two controllers.

The evaluation of DEAR in Abilene topology (11 nodes, 28 links) is shown in Figure 1 for the case of having two controller domains, against two other versions of the algorithm with additional constraints (that is MLU and Control Traffic Delay constraints). The controllers placement were obtained using the well known minimum k-median model (Heller et al., 2012).

Results show that DEAR could save until near to 40% of energy consumption when traffic is low. It is also shown, that more restrictive constraints will be paid with less energy saving. This behavior is expected given that, in order to meet the new performance requirements, a fewer number of alternate paths can be considered in the optimization. Therefore, it will be a trade-off to consider in accordance with the main objectives in each implementation.

In a second set of simulations, the analysis of using a modified Shortest Path Routing (Mod-SPR) is also included to get a sense of the values of energy savings achieved by our approach. Mod-SPR can be considered as a default shortest path routing algorithm for multiple domains SDN with in-band control traffic, where data plane communications cannot be routed through any controller. We use Mod-SPR as a fair comparison in the evaluation since there is no research considering energy saving with in-band control traffic in multiple domains SDN under routing behaviour presented in this proposal.




Figure 2: DEAR vs. Mod-SPR in Abilene topology with two controllers.

As shown in Figure 2 in all cases our distributed energy-aware routing approach outperforms the shortest path routing in terms of energy saving. In general DEAR achieves significant energy savings but bigger improvements over shortest path routing are reached when the traffic grows.

4. Conclusions

In this paper, we proposed a distributed energy-aware routing approach that optimizes the number of active links required to route the control and data plane communications in large-scale SDN with multiple domains. Such goal is achieved by an ILP model with two steps of optimization that integrates the routing requirements for data and control traffic in OpenFlow networks with in-band control traffic.

Using an agent-based approach, DEAR can be implemented as a software agent in each one of the distributed controllers in different SDN domains. In this way, an energy-aware control plane could be achieved, where the controllers determine the link interfaces that should be put into sleep mode. This proposal allows to attain optimal solutions for the power consumption problem in a multi-domain SDN. Based on experimental simulations using a real topology and traffic demands, we showed that our distributed energy-aware routing approach achieves energy savings of up to 40% and outperforms the shortest path routing with noticeable improvements.

Developing a heuristic algorithm to use this model in topologies with a bigger number of nodes in each domain, will be an important task as future work. We also plan to extend this work to take into account the use of restoration mechanisms in order to improve the fault tolerance capacity of our approach.

5. Acknowledgments

This work has been supported by the Ministerio de Economía y Competitividad of the Spanish Government under project TEC2013-47960-C4-1-P and through a predoctoral FPI scholarship.

6. References

- Aebischer, B. and Hilty, L. M., 2015. The Energy Demand of ICT: A Historical Perspective and Current Methodological Challenges. In *ICT Innovations for Sustainability*, chapter 4, pages 71–103. Springer.
- Gelenbe, E. and Caseau, Y., 2015. The Impact of Information Technology on Energy Consumption and Carbon Emissions. *ACM Ubiquity*, 2015(June):1–15.
- Giroire, F., Moulrierac, J., and Phan, T. K., 2014. Optimizing Rule Placement in Software-Defined Networks for Energy-Aware Routing. In *Proc. IEEE GLOBECOM'14*, pages 2523–2529.
- Gupta, M. and Singh, S., 2003. Greening of the Internet. In *Proc. ACM SIGCOMM'03*, pages 19–26.
- Gurobi Optimization. Version 6.5. <http://www.gurobi.com/>. Last accessed on May 16, 2016.
- Heller, B., Sherwood, R., and McKeown, N., 2012. The Controller Placement Problem. In *Proc. HotSDN'12*, pages 7–12.
- Kreutz, D., Ramos, F. M. V., Veríssimo, P. E., Rothenberg, C. E., Azodolmolky, S., and Uhlig, S., 2015. Software-Defined Networking: A Comprehensive Survey. *Proceedings of the IEEE*, 103:14–76.
- McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rexford, J., Shenker, S., and Turner, J., 2008. OpenFlow:Enabling Innovation in Campus Networks. *ACM SIGCOMM Computer Communication Review*, 38(2):69–74.
- Sharma, S., Staessens, D., Colle, D., Pickavet, M., and Demeester, P., 2013. Automatic Bootstrapping of OpenFlow Networks. In *Proc. IEEE LAN-MAN'13*, pages 1–6.
- The Climate Group, 2008. SMART 2020 Report, Enabling the Low Carbon Economy in the Information Age. <http://www.theclimategroup.org/>. Last accessed on March 14, 2016.
- Van Heddeghem, W., Lambert, S., Lannoo, B., Colle, D., Pickavet, M., and Demeester, P., 2014. Trends in Worldwide ICT Electricity Consumption from 2007 to 2012. *Computer Communications*, 50:64–76.
- Wang, H., Li, Y., Jin, D., Hui, P., and Wu, J., 2016. Saving Energy in Partially Deployed Software Defined Networks. *IEEE Transactions on Computers*, 65(5):1578–1592.
- Wang, R., Jiang, Z., Gao, S., Yang, W., Xia, Y., and Zhu, M., 2014. Energy-Aware Routing Algorithms in Software-Defined Networks. In *Proc. IEEE WoWMoM'14*, pages 1–6.
- Zhang, M., Yi, C., Liu, B., and Zhang, B., 2010. GreenTE: Power-Aware Traffic Engineering. In *Proc. IEEE ICNP'10*, pages 21–30.
- Zhang, Y., 2004. Abilene TM. <http://www.cs.utexas.edu/~yzhang/research/AbileneTM/>. Last accessed on January 20, 2016.

