Energy-Efficient Software Defined Networks

Supervision

Leaders:

Frédéric Giroire, CNRS Researcher, COATI team (*).
Dino Martin Lopez Pacheco, Assistant Professor, SIGNET team.

Guillaume Urvoy-Keller (Professor, SIGNET team) and Joanna Moulierac (Assistant Professor, COATI team) will also actively participate to this project.

(*) Frédéric Giroire is currently preparing his HDR and intends to benefit from the free token offered by the doctoral school to be the official supervisor. Note that Guillaume Urvoy-Keller can alternatively be the official supervisor.

Scientific themes

Energy Efficiency

Context

Software-defined or Software-Driven Networks (SDN) is a new networking paradigm enabling innovation, centralization of network management and preventing the so-called ossification of the Internet. SDN decouples the control plane from the data plane in network equipments, which means that a switch or a router is transformed into a simple forwarding device that applies rules sent by a remote controller using a normalized protocol. This simple approach allows network administrators to get a better control on the traffic in their network, e.g., Google has recently presented an SDN-based re-design of its core backbone where it is able to reach nearly 100% utilization of links under stringent QoS constraints [8]. SDN also enables the academic community to experiment with flexible as well as high performing equipments to test new or existing protocols. The OpenFlow protocol is the leading instantiation of the SDN concept at the moment and is supported by major manufacturers, e.g., HP, Juniper, IBM as well as open-source virtual switches like Open vSwitch [10], which is at the heart of cloud management solutions like OpenStack [11].

In parallel, energy consumption of networking infrastructures (telecom, data center, ISP or enterprise networks) is a growing concern. For instance, the networking appliances represent alone 10 to 20% of the total energy cost of data centers [12]. Several energy-efficient algorithms have been proposed and they imply in most cases to adapt resource usage to the dynamics of the demand, for example in energy aware routing [5,7], base-station assignment [6] or dynamic rate adaptation [17]. Legacy network protocols and equipments allow only with difficulties to adapt to demand dynamics. Indeed, the forwarding policies should be added manually in each router, meaning a hard work to be done by the network administrators.

The SDN technology has the potential of making such an adaptation possible, leading to substantial energy savings. Indeed, SDN enables:

- A reporting of node and link metrology data,
- A centralized management using this data, allowing to execute more efficient computations and network optimization than in a distributed control manner (like legacy routing algorithms do),
- A fast reshaping of the traffic per port through modifications of the flow table which is directly exposed to the SDN controller.

The latter property is the key to turn on/off or switch the rate of a network interface, which are the two approaches at the root of the techniques proposed to reduce energy consumption [9,13]. One thus sees that the SDN technologies provides a promising environment to build new and efficient energy-aware protocols. Our objective in this thesis is to tackle the theoretical and practical challenges that will lead to such protocols.

Objective

In this PhD, we want to focus on the design of new energy-aware protocols and their deployment in Software Defined Networks. Specifically, we envisage to address both algorithmic and practical challenges considering different types of networks including backbone, data center, infrastructure-based wireless access (Cellular/Wifi) or enterprise networks.

From an algorithmic point of view, the challenge will be to propose energy-efficient algorithms and protocols while having in mind the advantages as well as the limitations of the SDN technology. We will answer the new algorithmic challenges introduced by this paradigm:

- In the new SDN compatible hardware, the forwarding table requires to be implemented in Ternary Content Addressable Memory (TCAM) which is expensive, power-hungry and physically bigger than standard RAM [3]. Consequently, the size of TCAM in SDN-capable switches is small and the number of entries (flows) in the forwarding table is very limited (e.g., 750 entries for the NEC PF5820, or nearly 4000 entries on the MLX). We need new algorithmic solutions to mitigate this problem and to address this scalability issue [2,4]. We need also to partition all these rules among the set of SDN nodes (that can be hierarchically organised) in order to reduce the need to consult the controller when new flows arrive.
- The virtualization and cloud services should be proposed with a rapid provisioning for the users together with an efficient resource management and placement. The mapping between virtual and physical networks should be also efficient in terms of QoS and resources availability. Consolidation of virtual machines (i.e. strategic deployment of VMs for an efficient usage of network resources) is also a way to reduce energy consumption both at the server and network level. Indeed, designing and implementing an optimal algorithm in the Data Center Manager that first consolidates virtual machines to minimize the number of active servers and, in a second stage, the set of active links and forwarding devices to be turned-off (with an adjustable safety margin) is an interesting avenue to further reduce energy consumption of data centers as compared to a pure networking approach [14].
- In SDN, when the fields of an incoming packet do not match any of the available rules in

the forwarding device, such a device must store the packet in the buffer and contact the controller in order to obtain the forwarding rule to apply. Hence, we will need also to determine both the optimal number of SDN controllers and their location within the network considered, in order to reduce the latency induced by the communication between the controllers and the SDN nodes and reduce the single-point of failures problem.

From a practical viewpoint, we will focus on the implementation of our green protocols in SDN testbeds and the study of the impact that such green protocols will have on the Quality of Service. In order to do this, we intend to:

- Implement our green protocol in OpenFlow as independent modules to be added to the OpenFlow architecture, and provide the needed interfaces in order to support the heterogeneity of the network equipments and the evolution of the network. Indeed, networks are built using different switches and router models with different power consumption models.
- Evaluate the operational efficiency of our OpenFlow module implementations with testbeds. These testbeds will be based on emulation and experimentation. Emulation relies on virtualization techniques with high-end servers emulating servers, switches and routers. Mininet [16] offers for instance an emulation environment based on OpenFlow in order to provide virtual switches with Open vSwitches. Mininet was used in [9] to build a hybrid evaluation model of a data center with a few physical servers and switches coupled with a mininet network. We also intend to build a testbed with commercial OpenFlow switches [13]. Additionally, in order to test our proposals, realistic traffic workload will be devised out of publicly available traffic traces from data centers [15], which was used in [9,14] for instance. We will also benefit from traffic traces of MapReduce applications collected on a private data center used for research and teaching in the group of Pietro Michiardi at Eurecom. Last but not least, Orange Labs will give us access to traffic matrices of a typical ISP network.
- Evaluate the impact of energy-efficient mechanisms proposed by us or others on the applications flowing in the network (e.g. Video broadcasting, Online games, HTTP servers, etc). Indeed, re-routing existing flows can impact the end-to-end jitter and introduce reordering with negative effects on the customers' required QoS. Similar problems on the jitter can be produced if the TCAM table is full and different forwarding rules must be applied, since before forwarding a packet, an OpenFlow switch must first contact the OpenFlow controller in order to take a decision; thus delaying the traffic. This problem is often overlooked in works focusing on reducing energy consumption. In [9], the authors focus on the queuing delay and other application level metrics but only for the case of UDP traffic, leaving the TCP case as future work, while TCP is the dominant transport protocol in mobile networks, data centers, enterprise network and the public Internet itself.
- Provide a robust solution. In presence of low traffic load and after shutting down some network equipments or ports, an unexpected increase of the traffic load due to rare events (like cable cuts or celebrity news [13]) can lead to heavy congestions. In this project, we will evaluate the capacity of our propositions to handle this kind of events (e.g. how fast does it converge to the right network configuration). If needed, we will also propose a solution able to handle sudden variations in traffic.

Complementary of teams involved in the project

This project is supported by both COATI and SIGNET.

COATI (Combinatorics, Optimization, and Algorithms for TelecommunIcations) is a joint project-team between INRIA Sophia Antipolis – Méditerranée and the I3S laboratory, with a strong expertise in algorithms, graph theory and combinatorial optimization applied to telecommunication (wired and wireless) networks. This is a key expertise to tackle the algorithmic challenges of the proposal. The members of COATI involved in this project have worked extensively on devising energy efficient algorithms for backbone and access networks.

SIGNET (from SIGnal to NETworks) is a research team of I3S with expertise in networking and mobile communication networks. The SIGNET people involved in this work will bring their expertise in the simulation of large networks and in doing fine-grained analysis of the application/transport layer mechanisms, as well as the impact of virtualization.

Note that this PhD proposal would constitute a first collaboration between COATI and SIGNET.

Selected publications

[p1] Araujo, J., Giroire, F., Liu, Y., Modrzejewski, R., Moulierac, J. Energy efficient content distribution. In *IEEE ICC* 2013.

[p2] Modrzejewski M., Chiaraviglio, L., Issam, T., Giroire, F., Le Rouzic, E., Bonetto, E., Guerero, C.Energy Efficient Content Distribution in an ISP Network. In *IEEE GLOBECOM* 2013

[p3] Giroire, F., Moulierac, J., Phan, T. K., Roudaut, F. Minimization of network power consumption with redundancy elimination. In *NETWORKING* 2012.

[p4] Hafsaoui, A., Urvoy-Keller, G., Siekkinen, M., Collange, M.: Detecting and Profiling TCP Connections Experiencing Abnormal Performance. TMA 2012.

[p5] Orgerie, A.-C., Lefèvre, L., Guérin Lassous, I., López-Pacheco, D.M.: ECOFEN: An End-to-end energy Cost mOdel and simulator For Evaluating power consumption in large-scale Networks. *IEEE WOWMOM* 2011.

[p6] Arsene, A., Lopez-Pacheco, D.M., Urvoy-Keller, G.: Understanding the network level performance of virtualization solutions. *IEEE CLOUDNET* 2012.

References

[1] AGARWAL, S., KODIALAM, M. and LAKSHMAN. Traffic engineering in software defined networks. *In IEEE INFOCOM 2013*

[2] KANG, Nanxi, LIU, Zhenming, REXFORD, Jennifer, *et al.* Optimizing the “One Big Switch” Abstraction in Software-Defined Networks. *Proc. ACM CoNEXT*, 2013.

[3] STEPHENS, Brent, COX, Alan, FELTER, Wes, *et al.* PAST: Scalable Ethernet for data centers. In : *Proceedings of the 8th international conference on Emerging networking experiments and technologies*. ACM, 2012. p. 49-60.

[4] KANIZO, Yossi, HAY, David, et KESLASSY, Isaac. Palette: Distributing tables in software-defined networks. In : *IEEE Infocom Mini-conference*. 2013.

[5] Chiaraviglio, L., Mellia, M., & Neri, F. (2009, June). Reducing power consumption in backbone networks. In *Communications, 2009. ICC'09. IEEE International Conference on* (pp. 1-6). IEEE.

[6] MARSAN, M. Ajmone, CHIARAVIGLIO, Luca, CIULLO, Delia, *et al.* Optimal energy savings in cellular access

networks. In : *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*. IEEE, 2009. p. 1-5.
[7] GIROIRE, Frédéric, MOULIERAC, Joanna, PHAN, Truong Khoa, *et al.* Minimization of network power consumption with redundancy elimination. In : *NETWORKING 2012*. Springer Berlin Heidelberg, 2012. p. 247-258.
[8] JAIN, S. *et al*, B4: experience with a globally-deployed software defined wan. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM* (SIGCOMM '13). ACM, New York, NY, USA, 3-14.
[9] NGUYEN HUU, T. *et al.* Modeling and experimenting combined smart sleep and power scaling algorithms in energy-aware data center networks. Simulation Modelling Practice and Theory, Volume 39, December 2013, Pages 20–40
[10] http://openvswitch.org/
[11] http://www.openstack.org/
[12] GREENBER A., HAMILTON J., MALTZ D. A., and PATEL P. , "The cost of a cloud: research problems in data center networks," *SIGCOMM Comput. Commun. Rev.*, pp. 68–73, 2008.
[13] HELLER, Brandon ; SEETHARAMAN, Srinivasan ; MAHADEVAN, Priya ; YIAKOUMIS, Yiannis ; SHARMA Puneet ; BANERJEE, Sujata ; McKEOWN Nick: ElasticTree: Saving Energy in Data Center Networks.. In: *NSDI* : USENIX Association, 2010.
[14] Shirayanagi, Hiroki ; Yamada, Hiroshi ; Kono, Kenji: Honeyguide: A VM migration-aware network topology for saving energy consumption in data center networks.. In: ISCC : IEEE, 2012. - ISBN 978-1-4673-2712-1, S. 460-467
[15] http://pages.cs.wisc.edu/~tbenson/IMC10_Data.html
[16] http://mininet.org
[17] Nedevschi, S., Popa, L., Iannaccone, G., Ratnasamy, S., & Wetherall, D. Reducing Network Energy Consumption via Sleeping and Rate-Adaptation. In *NSDI* 2008.