Ecole Doctorale STIC

Sciences & Technologies de l'Information et de la Communication

# ED STIC - Proposition de Sujets de Thèse

## pour la campagne d'Allocation de thèses 2015

| | |
|---|---|
| **Axe Sophi@Stic :** | SystèmesRéseaux\| |
| **Titre du sujet :** | How sustainable data centers can be |
| **Mention de thèse :** | Informatique |
| **HDR Directeur de thèse inscrit à l'ED STIC :** | Philippe Nain |

**Co-encadrant de thèse éventuel :**

| | |
|---|---|
| **Nom :** | Alouf et Hermenier |
| **Prénom :** | Sara et Fabien |
| **Email :** | Sara.Alouf@inria.fr,Fabien.Hermenier@unice.fr |
| **Téléphone :** | 04.92.38.71.99 |

| | |
|---|---|
| **Email de contact pour ce sujet :** | Sara.Alouf@inria.fr,Fabien.Hermenier@unice.fr |
| **Laboratoire d'accueil :** | I3S/INRIA |

**Description du sujet :**

SUMMARY

The objective of this thesis is to investigate how sustainable data centers can be. The contributions of the thesis will be both theoretical and experimental, aiming to provide insights into the following questions: what are the key features of a green task scheduler? Is it possible to achieve net-zero* grid power in the absence of energy storage? Is it possible to use only renewable energy?

CONTEXT

Data centers are facilities that house a potentially huge amount of computer systems that are used mainly for storage or computing. Typically, there are three subsystems in a data center: the equipment, the power infrastructure and the cooling infrastructure. To reduce the ecological

footprint of data centers, one needs to optimize the three subsystems [4]. Better equipment may consume less power; using renewable energy reduces the electricity bill and improves sustainability; innovative solutions can improve the efficiency of the cooling infrastructure and reduce its power consumption. Workload management is one way of improving the sustainability of green data center [2, 3]. When the power infrastructure includes sources of renewable energy, it is highly interesting to shift as much of the load to the periods of high-energy production. This however may conflict with the best operation time seen by the cooling infrastructure. For instance, solar energy is abundant during clear days but at the same time, the temperature is at its highest point, which is the least convenient time to run servers at maximum load (the cooling infrastructure would need to be used at its maximum potential).

Another point to be taken into consideration is the type of services a data center offers as its energy needs may differ substantially from one data center to another. Here, we are not referring to the bulk of energy needed in the long run but rather to the short-term consumption as it depends on the characteristics of the jobs being served. Basically, service-oriented or batch-oriented jobs, bag of tasks or workflows, rigid or malleable jobs do not offer the same management opportunities and renewable energy sources may be enough to power a data center in some contexts but insufficient in others. State of the art techniques only focus on workload management solutions to control the data center consumption according to the energy availability [2, 3, 4]. We think it would also be meaningful to consider capacity planning techniques to infer the best data center architecture possible. For example, the usage of batteries for energy storage may be crucial in some regions to host particular workloads. Similarly, the number of photovoltaic arrays or wind turbines should be evaluated carefully to estimate the most efficient architecture according to a workload, the possible penalties when jobs are delayed and the economical environment.

THESIS' OBJECTIVES

Given the advisors' expertise, the focus will be on workload management techniques as a way to improve the sustainability of data centers. A first objective of the thesis is to implement a green scheduler. A second objective relates to the evaluation of the green scheduler. This evaluation will be performed analytically, using a stochastic model for the jobs and the renewable energy production, but also empirically. Traces of the jobs and of the energy consumption of two data centers (Grid5K and the one of the University of Nice Sophia Antipolis) will be used to (1) validate the model and to (2) test the green scheduler in different conditions. Ultimately, the contributions of the thesis will provide insights into the following questions: what are the key features of a green task scheduler? Is it possible to achieve net-zero grid power consumption in the absence of energy storage? What is the price to pay (economical, environmental, serviceability) to provide a data center only powered by renewable energy?

* Net-zero grid power consumption refers to the situation in which the overall energy consumed (renewable and/or from the grid) during a given period of time does not exceed the total renewable energy generated during the same period of time.

SOME KEY REFERENCES (NON-EXHAUSTIVE LIST)

[1] BtrPlace, An Open-Source flexible virtual machine scheduler, http://www.btrplace.org/.

[2] I. Goiri, W. Katsak, K. Le, T.D. Nguyen, R. Bianchini, “Parasol and GreenSwitch: managing datacenters powered by renewable energy,” Proc. of ASPLOS 2013, Pages 51-64.

[3] I. Goiri, K. Le, M.E. Haque, R. Beauchea, T.D. Nguyen, J. Guitart, J. Torres, R. Bianchini, “GreenSlot: Scheduling energy consumption in green datacenters,” Proc. of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Seattle, Washington, USA, 12-18 Nov. 2011, pp.1-11.

[4] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, C. Hyser, “Renewable and cooling aware workload management for sustainable data centers,” Proc. of ACM SIGMETRICS/PERFORMANCE 2012. (ACM SIGMETRICS Performance Evaluation Review, Volume 40 Issue 1, June 2012, Pages 175-186).

**English version:**