

DC4Cities: Better Usage of the Renewable Energies in Data Centres

Corentin Dupont
Create-Net
Trento, Italy
Email: cdupont@create-net.org

Fabien Hermenier
Univ. Nice Sophia Antipolis
CNRS, I3S, UMR 7271,
France
Email: fabien.hermenier@unice.fr

Data centres are a key part of the ICT chain. The energy consumed by them in 2010 accounted for between 1.1% and 1.5% of the total electricity consumed worldwide [1], making of data centre energy management an important challenge for researchers. With the recent adoption of renewable energies to power data centres [2], the research community enlarges its vision to associate with purely *quantitative* energy consumption reduction, the notion of *quality* of the energy consumed, *i.e.* the capacity to rely as much as possible on sustainable power sources. Differently from non-renewable energy sources, the availability of renewable energies is very volatile and time dependent: *e.g.* solar power is obtainable only during the day, and is subject to variations due to the meteorological conditions. The goal in this case is to shift the workload of running applications, according to the forecasted availability.




Figure 1. Adapting applications for a better usage of renewable energies

In our scenario, an entity external to the data centre called the Energy Management Authority is able to set high-level energy and power objectives for the data centre. For example, an objective can be: “at least 80% of the data centre energy consumption has to be generated from renewable energy sources”. To translate these high-level objectives into concrete actions performed in the data centre, our platform receives forecasts regarding the amount of available power and the energy source mix of its electricity inputs. These forecasts are either provided directly by the power suppliers or created by algorithms which use environmental and historical data to predict future renewable energy generation. Having both workload estimations for the different running applications [3] and the energy forecasts, our platform is able to assign specific power budgets to the different applications.

In order to run the applications in a data centre while respecting this power budget, we introduce the concept of Energy Aware Software Controller (EASC). The EASC is able to: (i) manage elasticity and scalability of multi-tier applications according to their energy footprint (*energy-awareness*); (ii) inject mechanisms in the single applications to negotiate energy

usage according to energy constraints (*energy-adaptiveness*).

The EASC implements workload scheduling techniques that can reconfigure the applications according to the power budgets (see Figure 1). An EASC is attached to each specific application running in the data centre and receives a different power budget, corresponding to the expected power consumed by the controlled application. The EASC is then proposing several plans to perform the application’s tasks, within the limits of its SLA. Those plans can be influenced using two weighted parameters: the aggressiveness and the eagerness. The aggressiveness controls the possibility for the application to consume more or less aggressively the renewable energies. For example, at a high aggressivity level the application will run at the highest performance level when renewable energies are available, and at the lowest performance when they are not. The eagerness controls the necessity for an application to complete its tasks more or less early. A low eagerness allow to be more flexible with regard to the scheduling of the applications tasks during the period of availability of the renewable energies. The EASCs are coordinated by the central component of our platform prototype. This central component coordinates the EASCs and selects the correct plan for each EASC, in order to optimize their combination.

A specific EASC has been designed to control the IaaS layer of modern data centres. It is used to consolidate the VMs that are running the applications, in order to free up some servers and switch them off when there is a low percentage of renewable energies. Similarly, an EASC for the PaaS layer is used to scale up and down PaaS applications with regard to their tasks, SLA and energetic situation.

Our solution has been trialled in the data centre of the health agency of the province of Trento, Italy. In this trial, the EASC was controlling a real application producing medical reports. We used an SLA stating that the application should be able to create 780 reports per day at minimum. Using our prototype, the percentage of renewable energy consumed by the data centre (called REN%) went from 43.1% up to 57.9%.

I. ACKNOWLEDGMENTS

The authors would like to thank the EU FP7 project DC4Cities (grant agreement number 609304), and the consortium members.

REFERENCES

- [1] Jonathan Koomey. Growth in data center electricity use 2005 to 2010. Technical report, Analytics Press, July 2011.
- [2] Íñigo Goiri, William Katsak, Kien Le, Thu D. Nguyen, and Ricardo Bianchini. Parasol and greenswitch: Managing datacenters powered by renewable energy. *SIGPLAN Not.*, 48(4):51–64, March 2013.
- [3] Aman Kansal and Feng Zhao. Fine-grained energy profiling for power-aware application design. *SIGMETRICS Perform. Eval. Rev.*, 36(2):26–31, August 2008.