Future Generation Computer Systems 55 (2016) 480–494

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Estimation and forecasting of ecological efficiency of virtual machines

Gregory Katsaros [c,*], Pascal Stichler [c], Josep Subirats [a,b], Jordi Guitart [b,a]

[a] Barcelona Supercomputing Centre (BSC), 08034 Barcelona, Spain
[b] Universitat Politecnica de Catalunya - Barcelona Tech (UPC), 08034 Barcelona, Spain
[c] FZI - Research Center for Information Technology, Berlin, Germany

HIGHLIGHTS

- PoC which demonstrates the eco-efficiency monitoring and its impact on VM placement.
- Monitoring and assessment framework and data structure is being designed.
- Modeling and definition of Ecological Efficiency of cloud VMs.
- Prediction of eco-efficiency and assessment using Correlation Coefficient.
- Extensive results from deployments onto two testbeds within Germany regions.

ARTICLE INFO

Article history:
Received 21 April 2014
Received in revised form
8 December 2014
Accepted 12 January 2015
Available online 20 January 2015

Keywords:
Ecological efficiency
Cloud computing
Virtual machine
Energy consumption
Monitoring

ABSTRACT

The massive development of the cloud marketplace is leading to an increase in the number of the Data Centers (DCs) globally and eventually to an increase of the CO_2 related footprint. The calculation of the impact of Virtual Machines (VMs) on the environment is a challenging task, not only due to the technical difficulties but also due to the lack of information from the energy providers. The ecological efficiency of a system captures the relationship between the performance of the system with its environmental footprint. In this paper we present a methodology for the estimation and prediction of the ecological efficiency of VMs in private cloud infrastructures. We specifically focus on the information management starting from the energy resources in a region, the energy consumption and the performance of the resources and finally the calculation of ecological efficiency of a VM. To this end, we have designed and implemented a framework through which the ecological efficiency of a running VM can be assessed and the ecological efficiency of a VM to be deployed can be forecasted. The presented framework is being evaluated through several private cloud scenarios with VM deployments in hosts located in Germany.

1. Introduction

The rapid growth of ICT application services goes along with an increase in number and size of data centers (DCs) that host these services. Because data centers are massive energy consumers, the carbon footprint of application services is moving more and more into focus. It is considered that ICT presently accounts for approximately 2% of global carbon emissions (more than 830 million tons of carbon dioxide) [1,2]. The advent of the cloud computing paradigm gave an enormous boost to the ICT services sector and the prediction towards 2020 is that the market (SaaS, PaaS, IaaS) will quadruple from what it is today.

NESSI, in the Strategic Research and Innovation Agenda for 2013 [3], clearly defined research priorities and recommendations: we must find new ways to increase software performance and energy-efficiency, by engineering energy-aware software to improve power-efficiency of software systems and services. Apart from the pure environmental and technological perspective, one should consider the current market demands. There is a significant part of the consumers community who enjoys above-average income and spending, and their buying decisions reflect environmental considerations [4]. The market of green consumers has an estimated global potential of 200 billion to 400 billion. Approximately 10% of German consumers and 12% of US consumers are willing to pay 10% more for "green" products that require less energy to operate or are manufactured by companies with a "green" reputation. A recent report [5] also shows that 28% of IT decision-makers at companies in Europe, the Middle East and Africa consider environmental criteria to be very important when

* Corresponding author.
E-mail addresses: katsaros@fzi.de (G. Katsaros), stichler@fzi.de (P. Stichler), josep.subirats@bsc.es (J. Subirats), jordi.guitart@bsc.es (J. Guitart).

http://dx.doi.org/10.1016/j.future.2015.01.002
0167-739X/

purchasing IT products and services. Another 67% of those surveyed attach relatively high importance to environmental criteria.

Until now the efforts of Green ICT in relation with cloud computing were restricted at the level of the Data Center and mainly in the form of increasing the energy efficiency of the infrastructures (PUE). While this activity represents only a small part of the cloud marketplace, mainly among the IaaS providers, the challenging question is whether we can introduce ecologically friendly policies and concepts in the level of applications and service providers. In that context, the element that is the main building block of all modern, cloud-enabled applications is the Virtual Machine (VM). Thus, the starting point of a relevant study and investigation should be at the level of virtual infrastructure.

Ecological efficiency of a system is the amount of work that is delivered in relation to its CO_2 emissions. To assess the eco-efficiency of cloud-based services, it becomes increasingly important to investigate the eco-efficiency of the cloud resources that the service utilizes. To this end, there are several technical constraints and challenges [6]: virtualization technology, which is a major characteristic of cloud computing, introduces an abstraction layer between the consumers of cloud resources and the physical infrastructure. The energy consumed by a service or a VM cannot be directly metered and therefore must be estimated through certain modeling methodology. What is more, the CO_2 footprint of the cloud resources is directly related with the energy-mix that the respective data center consumes at the time of the VM operation. Considering that the needed energy is provided by the local power providers, the calculation of ecological efficiency in cloud computing is therefore a location and time relative figure.

The above-mentioned issues harden the effective calculation of ecological efficiency of cloud services. To ecologically evaluate the resources that a cloud-enabled application utilizes, providers require appropriate tools and methods that are still not there. Thus in this paper we aim at investigating the environmental impact of cloud resources and specifically estimate the ecological efficiency at the granularity of the VM. In particular, we propose and implement a methodology that allows assessing the ecological efficiency of running VMs in the provider and forecast the ecological efficiency of VMs to be deployed or in operation. This work constitutes an updated and significantly extended version of the concepts presented in [7].

As has been presented also in [8] there are many parameters that affect the cost (energy or financial) of cloud resources. We, though, selected to focus on the ecological efficiency of the VM as it is one very interesting variable parameter that directly relates performance with power consumption of cloud infrastructures, resulting in a highly dynamic metric. In our work, we have not considered the cost related with the static infrastructure (cooling of Data Center, management cost, etc.), because these figures can be calculated separately and added on top of the parameters that we calculate dynamically. In addition, we assume that contemporary cloud application topologies consist of multiple application components installed in different VMs. Therefore, the power consumption and eventually the ecological efficiency should be investigated at the level of VMs.

In the following sections we present the proposed methodology and corresponding implementation to calculate and forecast the ecological efficiency of a VM considering the location and time constraints of its operation. In Section 2 we present the state of the art and the related work in the respective fields of research. In Section 3 we elaborate on the proposed solution: a monitoring system that allows assessing and forecasting the eco-efficiency of VMs. We describe in detail the architecture of the solution as well as the theoretical methodology and concept. Finally, in Section 4 we proceed with the evaluation of the implemented solution with several deployment scenarios within Germany, while in Section 5 we conclude and summarize our findings.

2. State of the art and related work

In order to reduce the carbon footprint of a system or a service, first of all we need to monitor and analyze the performance as well as the energy related information of our computing infrastructure. In cloud computing, it is important to know how much energy a specific service or VM consumes, rather than the consumption of the whole physical infrastructure. However, measuring the energy consumption of a single or even several VMs is a challenging task. From a consumer point of view, VMs are a black box whose energy consumption can only be estimated [9]. In order to do so, power usage models are normally used where performance characteristics are being used for the modeling of the energy consumption. In [10–13] it is pointed out that CPU utilization is the main factor driving energy consumption of a computing system, with memory and disk resource utilization playing a secondary role.

What is more, in [14] they make use of power usage metering to calculate and forecast the energy efficiency level of VMs in order to optimize VM deployments in private clouds. The methodology for the calculation of the energy efficiency of all cloud entities has been proposed using the CPU utilization of the VMs as the parameter to define the useful work performed and linear regression technique for the forecasting. To this end, there have been many discussions about measuring the useful work of a computer accurately [15]. Several benchmarks have been proposed [16], however, since every application has different requirements no universal formula can be determined. In the same context, power consumption modeling has been used for power-aware VM allocation using genetic algorithms [17] or through heuristic algorithms [18]. The definition of energy efficiency in cloud computing has been playing an important role also for the application of VM consolidation strategies [19,20].

The term ecological efficiency is a rather generic description of the efficient use of ecological resources. In computing systems, eco-efficiency can be seen as computing power delivered compared to the environmental resources needed to do so. And again, in this context the complete product or service life-cycle and its related impacts on the environment have to be considered. In the context of cloud computing, there have been different approaches to define eco-efficiency. For example, Google publishes the carbon emissions per query (0.2 g CO_2) and per watched minute on YouTube (0.1 g of CO_2) [21]. Similarly, in [22] a charge-back model is presented, where the environmental impact of providing data center services to the service consumers is traced back to the consumer. The consumer receives information about the CO_2 intensity of each transaction as well as the overall CO_2 emissions produced by his transactions. The eco-efficiency is calculated by CO_2 emissions per data transaction executed on the service.

In [23], the authors have formulated the cost of VM migrations between private clouds aiming at the reduction of the carbon footprint of a cloud network. Furthermore, in [24] a routing methodology for user placement in data centers is presented, that generates minimum carbon footprint and therefore optimizes the ecological efficiency. Both research studies point out the significance of the geographic location of the cloud infrastructure towards the increasing of the eco-efficiency of a service deployment. In addition, in [25], a framework for optimizing the carbon efficiency in clouds is presented, which is based on the installation of a registry with offers from cloud providers including data about their CO_2 emission rate, average DCiE (Data Center infrastructure Efficiency), VM power efficiency, prices, etc.; all this information has to be updated by the provider though.

While the benchmarks described in [26] are used to measure the offline energy efficiency of a node or set of nodes, the algorithm described in [27] considers real-time metrics, but not the energy efficiency itself. It actually considers metrics which potentially affect the energy efficiency. In this publication, we evolve and