



Assessing and forecasting energy efficiency on Cloud computing platforms



Josep Subirats*, Jordi Guitart

Barcelona Supercomputing Centre (BSC), Jordi Girona 31, 08034 Barcelona, Spain
 Universitat Politècnica de Catalunya (UPC), Jordi Girona 1-3, 08034 Barcelona, Spain

HIGHLIGHTS

- Assess and forecast energy/ecological efficiency for multiple levels in real time.
- Assess and forecast energy/ecological efficiency for potential actions.
- Estimate the future CPU utilisation of a VM.

ARTICLE INFO

Article history:

Received 18 October 2013
 Received in revised form
 10 November 2014
 Accepted 10 November 2014
 Available online 18 November 2014

Keywords:

Cloud computing
 Energy efficiency
 Ecological efficiency
 Forecasting
 Green computing
 IaaS provider

ABSTRACT

IaaS providers have become interested in optimising their infrastructure energy efficiency. To do so, their VM placement algorithms need to know the current and future energy efficiency at different levels (Virtual Machine, node, infrastructure and service levels) and for potential actions such as service deployment or VM deployment, migration or cancellation. This publication provides a mathematical formulation for the previous aspects, as well as the design of a CPU utilisation estimator used to calculate the aforementioned forecasts. The correct adjustment of the estimators' configuration parameters has been proved to lead to considerable precision improvements. When running Web workloads, estimators focused on noise filtering provide the best precision even if they react slowly to changes, whereas reactive predictors are desirable for batch workloads. Furthermore, the precision when running batch workloads partially depends on each execution. Finally, it has been observed that the forecasts precision degradation as such forecasts are performed for a longer time period in the future is smaller when running web workloads.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Over the last years, Cloud computing has become a consolidated computing paradigm that allows sharing different kinds of resources over the network in an automated way. Infrastructure as a Service (IaaS) providers provide low level computing resources on which the user can run its own software, typically in the form of Virtual Machines (VMs) and Virtual Local Area Networks (VLANs) established between them. The use of VMs allows the creation of services that can ask for more or less resources depending on their instantaneous demand variation (scalability). For instance, by simply adding a new VM replica to an already existing service, it can process the requests of a bigger number of users of that service. Another advantage of using VMs is that they offer the possibility to

be migrated from one physical host to another without stopping its execution, transparently to the user.

Back in 2006, the annual energy consumption of the datacenters in the United States was already of about 61 Billion kWh, which is equivalent to the energy consumed by 5.8 million average US households [1]. From that year, this amount of energy has significantly increased. What is more, given that most of the energy produced in US and around the world comes from burning coal and natural gas [2], the datacenter's energy consumption has a direct impact on climate change [3].

An individual IaaS provider typically owns a very large number of physical hosts, where the different users' VMs are deployed. Altogether, they consume a massive amount of energy which represents about 30% of their datacenter's operating expenses [4]. Moreover, IaaS providers have to pay fines if their carbon footprint exceeds the limits imposed by international regulations such as the Kyoto Protocol [5]. As a consequence, IaaS providers have a great interest in reducing the total power consumption of their physical hosts.

Each server has a power consumption that is not proportional to its resource utilisation. Even though the power consumption

* Corresponding author at: Barcelona Supercomputing Centre (BSC), Jordi Girona 31, 08034 Barcelona, Spain.

E-mail addresses: josep.subirats@bsc.es (J. Subirats), jordi.guitart@bsc.es (J. Guitart).

grows linearly with regards to its utilisation [6], if the host is not used it still has an associated idle power consumption. Using VM migration, multiple VMs can be consolidated into a single physical server and the remaining idle servers can be turned off, avoiding their idle power consumption and therefore reducing the overall datacenter's energy consumption.

Until recently, IaaS providers only cared about providing the required performance to their users, regardless of the incurred energy consumption. Green Cloud computing is envisioned to achieve efficient processing while minimising the energy consumption [7]. Energy efficiency is defined as the ratio between the useful output of a process and the energy input to this process [8]. It is a generic term which can be applied in multiple fields, but when talking about Green Cloud computing, improving the energy efficiency refers to using less energy to produce the same amount of services or useful output. Therefore, the ultimate target of Green Cloud computing is maximising the energy efficiency.

IaaS providers need tools to measure and predict the energy efficiency of their infrastructure in real time (this includes the performance and the power being delivered), as well as the generated carbon emissions. The current available tools are in fact benchmarks which are able to evaluate the energy efficiency of a node or set of nodes in offline mode, but they are not able to evaluate it in real time while executing general applications, or to predict it for a given amount of time in the future. Some publications use real-time metrics affecting the energy efficiency to drive their VM placement algorithms, but not the energy efficiency itself. What is more, IaaS providers use VM placement algorithms which need to know what will be the impact of performing a VM deployment/migration/cancellation on the infrastructure's energy efficiency in order to optimise it. Even though some of these calculations have been addressed by the related literature, they have been only evaluated in simulated environments, but no tools have been provided to IaaS providers to gather such information.

In this context and taking into account the limitations of the current available tools, the present publication tackles the following innovations:

- Develop mechanisms to assess and forecast the energy efficiency and the ecological efficiency of a Cloud IaaS provider's infrastructure in real time. Different levels will be considered: infrastructure level, node level, service level and VM level.
- Develop mechanisms to predict the energy and ecological efficiency of a Cloud IaaS provider infrastructure for potential actions, including: service deployment, VM deployment, VM migration and VM cancellation.
- Develop a mechanism to estimate the future CPU utilisation of a VM.

The present publication is structured as follows. Section 2 provides the necessary background and exposes the related work. Section 3 contains all the mathematical models developed to evaluate and forecast the energy efficiency at different levels for the current status of the provider, as well as to forecast it for potential actions. In Section 4, the description of the CPU estimator indirectly used when predicting the energy efficiency can be found. The developed tool is evaluated in Section 5 using two kinds of realistic workload: web and batch. The configuration parameters of the CPU estimator are also tuned in order to obtain a good precision when forecasting the energy efficiency of the infrastructure. Finally, in Section 6, the conclusions drawn from this work and future research lines are presented.

2. Background and state of art

2.1. Background

Along with the outburst of Cloud computing came a surge of interest in Green Computing. The latter is a field where the engi-

neers and developers try to optimise the energy efficiency of various computing processes either by intervening into the hardware or the software architecture.

A tool to measure the energy and ecological efficiency (see Section 3) has been developed as a component called the Ecoefficiency Tool, as part of the OPTIMIS project [9] (Optimized Infrastructure Services, FP7-ICT-2009-5 program). A preliminary version was introduced in [10] and has been extended and improved in this publication. If a service (a set of VMs performing a collaborative task) is accepted by an IaaS provider, an OPTIMIS component called Virtual Machine Manager (VMM) evaluates what is the best node to deploy each of the VMs of the service, considering the Business Level Objective (BLO) of the provider. If the BLO seeks to maximise the energy/ecological efficiency of the infrastructure, the VMM will choose the destination node that maximises the infrastructure's energy/ecological efficiency if the VM is deployed in it. When the node to host a particular VM has been chosen, the VMM issues the VM creation by sending its VM description to the virtualisation middleware.

The Ecoefficiency Tool continuously monitors the current and future energy/ecological efficiency of the infrastructure, and notifies higher-level components if a potential threshold violation is expected to happen, so corrective actions (such as VM migration) can be taken in advance. The performed predictions need to be as accurate as possible in order to avoid unnecessary and incorrect notifications. This is the reason why the conducted experimentation has been focused on trying to maximise the precision of the tool when performing these forecasts. Additional actions like VM deployment or cancellation can also be considered for service scalability. Its impact in the infrastructure's energy/ecological efficiency can be calculated using the methods described in Section 3.2.

The Ecoefficiency Tool calculates/predicts the energy/ecological efficiency using three kinds of data types: performance, CPU utilisation and power data. We have adopted the MWIPS metric to measure the performance, which refers to the Millions of Whetstone Instructions per Second performed with respect to those performed by a SPARCstation 20-61 with 128 MB RAM, a SPARC Storage Array, and Solaris 2.3. The MWIPS of each physical node have been measured using the UnixBench benchmark [11]. The CPU utilisation data is captured in real time using the *virt-top* utility [12]. The power data is obtained in real time from a power sensor using the SNMP protocol.

The PUE or Power Usage Effectiveness is defined in [13] as the ratio between the total power consumed in a datacenter (P_{DC}) and the power delivered to the computing infrastructure (P_{IT}). As exposed in [14], P_{IT} actually depends, among others, on the temperature of the cold air supplied by the air conditioning to the infrastructure room (T_5), which also influences the PUE. For a fixed T_5 , the total power consumption of the datacenter grows linearly by a factor of PUE with regards to the IT power consumption P_{IT} . Even though T_5 could be modified and thus the PUE would vary, our assessments and forecasts are made for much smaller periods than the ones in which T_5 varies. In addition, as exposed in [13], different datacenters normally have different PUE values. However, we address the calculation of the energy efficiency at different levels for a single datacenter of an IaaS provider. For these reasons, the PUE can be considered as a constant in all the formulas.

The mathematical models described in this paper will be validated with different workload types. In particular, we differentiate between interactive and batch workloads:

- Interactive workloads have a variable amount of work to perform, which consists of a number of client transactions to process in a given instant. Therefore, resource utilisation for interactive workloads is highly variable over time, as it depends on the number of simultaneous clients interacting with the sys-

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات