Future Generation Computer Systems 32 (2014) 128–137

Contents lists available at ScienceDirect

Future Generation Computer Systems

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

Energy-credit scheduler: An energy-aware virtual machine scheduler for cloud systems

Nakku Kim [a], Jungwook Cho [b], Euiseong Seo [b,*]

[a] School of ECE, Ulsan National Institute of Science and Technology, Ulsan, 689-798, Republic of Korea

[b] College of ICE, Sungkyunkwan University, Suwon, 440-746, Republic of Korea

ARTICLE INFO

Article history:
Received 25 October 2011
Received in revised form
3 April 2012
Accepted 17 May 2012
Available online 28 May 2012

Keywords:
Virtualization
Scheduling
Energy-aware computing
Cloud computing
Resource accounting

ABSTRACT

Virtualization facilitates the provision of flexible resources and improves energy efficiency through the consolidation of virtualized servers into a smaller number of physical servers. As an increasingly essential component of the emerging cloud computing model, virtualized environments bill their users based on processor time or the number of virtual machine instances. However, accounting based only on the depreciation of server hardware is not sufficient because the cooling and energy costs for data centers will exceed the purchase costs for hardware. This paper suggests a model for estimating the energy consumption of each virtual machine without dedicated measurement hardware. Our model estimates the energy consumption of a virtual machine based on in-processor events generated by the virtual machine. Based on this estimation model, we also propose a virtual machine scheduling algorithm that can provide computing resources according to the energy budget of each virtual machine. The suggested schemes are implemented in the Xen virtualization system, and an evaluation shows that the suggested schemes estimate and provide energy consumption with errors of less than 5% of the total energy consumption.

1. Introduction

The energy costs of data centers are rapidly increasing. By 2012, these costs will be double the 2007 costs and will exceed the purchase costs for server hardware [1]. In addition to energy costs, cooling costs for data centers continue to increase because the ever-increasing power consumption by servers leads to increased heat dissipation. Thus, the cloud computing model is expected to become the next-generation computing infrastructure owing to its energy efficiency.

In a cloud system, virtualization is an essential tool for providing resources flexibly to each user and isolating security and stability issues from other users [2]. Because users' services and applications run on separate virtual machines (VMs), virtualization helps cloud platforms to accurately monitor and control the amount of resources being provided to users. Virtualization also facilitates the guarantee of service-level agreements (SLAs) between users and cloud platform providers by resolving QoS crosstalk between services and applications. Moreover, separation of the execution context and data between VMs is supported by processors and hypervisors so that users use their VMs as if they own their dedicated server hardware. Owing to these benefits, many commercial cloud systems including Amazon EC2 employ virtualization so that users can freely configure their virtual servers from kernel to application layers.

Currently, cloud services bill their users for the amount of computing resources provided [3]. In general, the amount of processor time and number of VMs allocated are common criteria for accounting [2]. Such billing systems are based on the common belief that the costs for purchasing and maintaining server hardware are proportional to the operation time. However, if significant factors that affect the management cost of servers are not considered, these billing schemes may have room for improvement. The energy consumption of a server system closely relates to the processor working time. However, the processor time does not accurately reflect the system energy consumption [4].

As stated earlier, energy and cooling costs will exceed hardware purchasing costs. Therefore, to fairly delegate energy costs to users, a billing system must infer the energy consumption for each user. However, measurement of the energy consumption for each VM running on a server with multiple VMs is a technically challenging problem. The use of embedded power meters in server systems is increasingly popular. However, because measurement of power consumption for the whole system is a totally different problem from measurement of that for only one VM, the problem remains even for embedded power meters.

* Corresponding author. Tel.: +82 31 299 4970; fax: +82 31 299 4921.
E-mail addresses: nkkim@unist.ac.kr (N. Kim), jmanbal@skku.edu (J. Cho), euiseong@gmail.com, euiseong@skku.edu (E. Seo).

0167-739X/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2012.05.019

Existing research has approached the measurement of per-VM energy consumption by accounting for and analyzing the resources used by each VM. However, most models for estimating energy consumption assume that the power consumption of each VM is simply proportional to the processor time allocated or the number of network packets processed.

In this paper, we reveal that processor time is not a satisfactory criterion for accurate estimation of power consumption by a VM and that some in-processor events affect processor power consumption more significantly than others. Based on observations made in diverse experiments, we suggest a model for estimating energy consumption that calculates the amount consumed by a VM via monitoring of processor performance counters.

Based on the proposed estimation scheme, we also propose an energy-aware VM scheduler that can limit the energy consumption of virtual machines to their energy budget. Conventional VM schedulers only consider the processor time when it comes to scheduling decisions. Different from its traditional counterparts, the energy-credit scheduler uses the energy consumption rates of VMs for scheduling. It schedules VMs so that their energy consumption rates remains below user-defined values.

Both the suggested estimation model and scheduler are implemented in the Xen virtual machine monitor. The implementation also provides external interfaces so that they can be easily adopted for real-world systems. We evaluate our prototype using the SPEC CPU2006 benchmark suite.

The remainder of the paper is organized as follows. The background and related works are introduced in Section 2. We suggest and evaluate a model for estimating the energy consumption of VMs in Section 3. Based on this model, we propose an energy-aware scheduler and evaluate it in Section 4. We conclude our research in Section 5.

2. Motivation and related work

2.1. Motivation

The amount of energy consumed by a server varies greatly depending on the workload it runs. Many studies on predicting or controlling server power consumption are based upon the assumption that the power consumption of a server is determined mainly by its processors, and that the power consumption of a processor is determined primarily by processor utilization [5,4,6]. However, energy consumption estimation based on processor utilization or execution time is generally not accurate enough for billing of energy costs.

Server vendors have recently begun to integrate power measurement hardware to monitor server power consumption on the fly [7]. By employing the integrated power measurement hardware, the energy consumption of a server can be measured accurately. However, an integrated power meter measures the power consumption of the whole server system, while the energy-based billing systems of cloud services require measurement of the energy consumption for each VM.

Despite previous research to leverage dedicated measurement devices, determination of the energy consumption of a VM is still a difficult problem for the following reasons.

First, the sampling intervals of integrated power meters generally fall between 1 and 0.1 s [8], whereas the time unit for scheduling of VMs is very short, ranging from a few hundred microseconds to a few tens of milliseconds. The time scale difference between measurement and scheduling makes it difficult to identify the energy consumption of a specific VM when there are multiple VMs competing for a processor.

Second, VMs in a multicore processor system share and compete with other VMs for system resources. Naturally, the throughput of a VM and its energy consumption may vary according to the characteristics and activities of other VMs that run concurrently [9]. Therefore, energy consumption measurement for a VM cannot rely on blind arithmetic that only considers processor and resource usage by the VM. Thus, to identify the energy consumption of each VM, which continually and rapidly changes, energy consumption estimates must consider both the activities of the VM and the effects from other VMs on resource utilization.

Besides the VM-level energy measurement schemes, energy-aware resource provisioning schemes are also essential for energy-cost-based billing systems so that users can limit the energy consumption of their VMs to stay within their energy budgets.

Most conventional VM schedulers aim to provide processor time to VMs fairly and proportionally according to their priorities because processor time is one of the major billing criteria of conventional computing systems. In other words, processor time is the major scheduling basis in conventional schedulers [10,11]. Therefore, a novel energy-aware scheduler and its premise must be designed from scratch. At the start of our research, we conceptualized a system for our aim as follows.

To utilize an energy-aware scheduling scheme, each VM should explicitly provide its energy budget to the scheduler. The energy budget of a VM is the amount of energy it is allowed to use during a *fiscal time interval*. If a VM uses up its energy budget in a fiscal interval, then that VM will be suspended until the current fiscal interval finishes and the next interval begins. The VM fiscal interval is determined by the owner of the VM according to its purpose and characteristics, and may range from a few tens of milliseconds to a few minutes. Thus, the fiscal interval of a VM may differ from that of other VMs.

The fiscal interval of a VM should be carefully chosen because it supposedly affects throughput and responsiveness of the VM. When the interval is too short, excessively frequent context switches will occur. If the interval is too long, the suspension time for VMs without any energy budget left will be too long and will negatively affect the responsiveness of the services provided by the VMs.

For example, as shown in Fig. 1, when a web service VM and a data backup service VM are run concurrently on a physical server, it is beneficial for the web service VM to have a short-term plan such as 300 mWh/0.5 s, while a long-term plan such as 18000 mWh/1 min would be more suitable for the backup service VM.

In this paper, we realize this conceptual model by suggesting a model for estimating per-VM energy consumption and an energy-aware VM scheduler.

2.2. Related work

There has been intensive research on power-capping resource management schemes, which limit the power consumption of a server or a set of servers to predefined thresholds [6,12]. Additional research has introduced techniques to distribute workloads and manage computing resources so that workloads are evenly distributed over a set of servers in terms of power requirements [13]. The goal of these efforts is to reduce the costs of building and maintaining a power distribution infrastructure, as well as energy costs, by reducing the peak power consumption of single nodes or server groups [4,14,13,12].

Predicting the power and energy consumption of a server that runs a specific workload is straightforward, especially when the workload is stable in terms of resource usage. However, predicting the power and energy consumption of a server on which multiple VMs are consolidated is difficult because the VMs interfere with each other. A statistical method has been proposed for predicting the power consumption of a server with