Computers and Electrical Engineering 47 (2015) 173–185

Contents lists available at ScienceDirect

Computers and Electrical Engineering

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

# Distributed consolidation of virtual machines for power efficiency in heterogeneous cloud data centers ☆

Hadi Khani [a], Amin Latifi [a], Nasser Yazdani [a], Siamak Mohammadi [b,*]

[a] Router Lab., School of Electrical and Computer Engineering, University of Tehran, Tehran, Iran
[b] Dependable System Design Lab., School of Electrical and Computer Engineering, University of Tehran, Tehran, Iran

ARTICLE INFO

*Article history:*
Received 15 October 2014
Received in revised form 1 August 2015
Accepted 3 August 2015
Available online 9 September 2015

*Keywords:*
Green computing
Virtualization
Game theory
Cloud computing
Data center

ABSTRACT

Data centers use dynamic virtual machine consolidation to reduce power consumption. Existing consolidation mechanisms are not efficient in cloud data centers, which have heterogeneous hardware infrastructure, huge scale, and highly variable non-stationary workloads. We use game theory to develop a novel distributed mechanism for both heterogeneous and homogeneous data centers of cloud computing. Our mathematical analysis shows that our mechanism converges after a finite number of migrations. In addition, we show that our worst case power consumption is only 23% more than the theoretical minimum. In order to validate our claim, we preform simulation in CloudSim with real workload traces from Google data centers.



## 1. Introduction

Power consumption in data centers has recently increased due to Cloud Computing era and greater demand for Big Data applications. Reducing power consumption has become a commitment for preserving the environment. Virtualization technology enables us to cut power consumption considerably. In this technology, multiple instances of operating systems share a server. The server is called physical machine (PM) and each operating system is called virtual machine (VM). Live migration is used to move running VMs between different PMs without disconnecting applications. The data centers can benefit form live migration to reduce the number of running PMs and turning off the idle ones.

In the past, different consolidation methods have been deployed by transforming the consolidation problem to the well-known bin packing problem. They use centralized algorithms, developed based on bin packing solutions, to assign VMs to PMs [1–4]. Some researchers have recently argued against the centralized algorithms and raised questions about their elasticity and scalability [5–8]. They have proposed distributed and self-organized approaches to tackle this issue. Unfortunately, simulation and lab-scale experiments are usually used to validate the proposed methods.

Several attempts have been made in order to use mathematics to design a more efficient distributed mechanism for this problem. Mastroianni et al. used mathematical analysis to tune their mechanism but they did not give any analysis about its performance [9]. Ye and Chen took a step forward and used mathematical analysis to establish a bound on the efficiency of their distributed algorithm [10]. However, they assumed that all PMs are identical. As the data center scale increases, the amount of heterogeneity in the data center also increases [11–13]. The assumption that all PMs have the same capacity is

☆ Reviews processed and recommended for publication to the Editor-in-Chief by Guest Editor Dr. R. Tolosana-Calasanz.

* Corresponding author.
*E-mail addresses:* hkhani@ut.ac.ir (H. Khani), a.latifi@ut.ac.ir (A. Latifi), yazdani@ut.ac.ir (N. Yazdani), smohamadi@ut.ac.ir (S. Mohammadi).

http://dx.doi.org/10.1016/j.compeleceng.2015.08.001
0045-7906/

obviously questionable [12,13]. Here, we develop a mathematical model for the dynamic consolidation problem in a data center equipped with heterogeneous servers.

### 1.1. Problem statement

We model a data center which consists of VMs hosted on PMs with different capacities. The PMs consume power proportionally to their loads. Our objective is to find the optimum mapping of VMs onto PMs that minimizes the social cost, i.e. the total power consumption of the whole data center. In order to meet agility and scalability, we assume that each VM can make decisions individually. In other words, there is no centralized decision maker.

Game theory has been used as a powerful tool to design and analyze some distributed systems and we believe that a cloud data center is not an exception. In our game model, we assume that each VM is the player that can choose any PM as a strategy. The chosen PM must certainly have sufficient residual capacity to host the VM. The player pays a cost depending on its strategy and the strategies of the others. The cost is the portion of power consumption of the chosen PM. Formally, when a VM with demand $d$ is hosted on a PM with power consumption $w$ and total demand $\varsigma$, its cost will be $w\frac{d}{\varsigma}$. Since all the players (VM) prefer to pay less; they favors PMs with the least residual capacity. This allows us to turn off unloaded PMs. Generally speaking, this is in line with our objective (minimum total power consumption), but needs more assessment since selfish behavior of VMs may keep final mapping away from the optimum.

In game theory, the first issue is whether the game converges to a stable point (Nash equilibrium). The game reaches a Nash equilibrium when no VM can benefit from changing its strategy unilaterally. The other issue is the inefficiency of power consumption in an equilibrium with respect to the optimum. We use the *price of anarchy* (PoA) and the *price of stability* (PoS) concepts to measure this inefficiency. The PoA is defined as the ratio between the social cost in the worst equilibrium and the optimum. Similarly, the PoS is the ratio between the social cost in the best equilibrium and the optimum. The PoA establishes an upper bound on the worst case result. The PoS has a more delicate interpretation. If $PoS = 1$ and we begin from an optimum mapping, then we can conclude that the selfish behavior of VMs cannot degrade that optimal mapping. Another important thing about the agility and applicability of the game in cloud computing is the convergence time. The convergences time is directly related to the number of migrations rquired to reach a Nash equilibrium.

### 1.2. Our contributions

The main contributions of this paper are:

- We propose a non-cooperative game for reducing power consumption in data centers with heterogeneous PMs.
- We show that the game always converges to a pure Nash equilibrium.
- We establish upper bounds on the convergence time of the game.
- We prove that $PoA = 1.23$. This means that a data center, using our method, consumes at most 23% more than its minimum required power.
- We show that $PoS > 1$. This means that the game sometimes degrades an optimal mapping and increases the total power consumption of the data center.

We evaluated our results by simulations using real-world workload traces from more than 6000 VMs running on a Google data center [14]. For our power model, we use the latest published results of the standard benchmarks [15]. Experimental results support the above findings and, more interestingly, our method reaches the optimal mapping significantly closer than indicated by the theoretical upper bound. In addition, the convergence time is many orders of magnitude shorter.

### 1.3. Outline of the paper

The rest of the paper is organized as follows. In the next section, we review the related work. In Section 3, first, we define our problem formally, and then present our mathematical analysis in depth. Section 4 is devoted to simulation results. We conclude in Section 5.

## 2. Related work

Dynamic virtual machine consolidation mechanisms can be broadly categorized as the follows:

- Heuristics based on algorithms for bin packing.
- Hierarchical solutions.
- Mathematical optimization approaches.
- Distributed and self-organized mechanisms.