Ain Shams University
Faculty of Computer and Information Sciences
Computer Science Department

Developing a Live Migration Aware Power Saving Algorithm in Cloud Computing

Thesis submitted as a partial fulfillment of the requirements for the degree of
Master of Science in Computer and Information Sciences

By
Salma Khaled Saad Elbay

Demonstrator at Computer Science Department,
Faculty of Computer and Information Sciences,
Ain Shams University

Under Supervision of

Prof. El-Sayed M. El-Horbaty
Professor of Computer Science

Dr. Islam Hegazy
Assistant Professor in Computer
Science Department

Faculty of Computer and Information Sciences,
Ain Shams University

Cairo – 2018

Acknowledgements

At this moment of accomplishment, All Praise to ALLAH, for giving me the blessing, the strength, the chance and endurance to complete this study.

I would like to express my gratitute to my thesis supervisor, Prof. El-sayed M. El-Horbaty for his guidance, motivation and insightful comments. Special mention goes to my enthusiastic supervisor, Dr. Islam Hegazy, who has shown me, by his example, how a good advisor should be. His guidance and constructive criticism helped me in all the time of research and writing of this thesis.

Nobody has been more important to me in the pursuit of this thesis than the members of my family. I would like to thank my parents; whose love and guidance are with me in whatever I pursue. They are the ultimate role models. Most importantly, I thank my loving and supportive husband, Amr, and my wonderful son, Yahia, whose existence provides unending inspiration. . . .

Abstract

The wide adoption of cloud computing is leading to an increase in the number of data centers and physical servers worldwide. Ecological communities are calling for the data centers to "go green" and hence, energy efficiency has become a crucial concern in modern data centers. Dynamic virtual machine (VM) consolidation is one of the effective approaches endorsed to achieve energy efficiency in cloud data centers hosting thousands of servers. Live migration is a core feature enabling virtual machine consolidation. However, live migration is a costly operation imposing energy and performance overhead. Thereby, an efficient dynamic virtual machine consolidate should consider the cost due to live migration.

This thesis addresses this challenge by presenting a dynamic virtual machine consolidation algorithm that is live migration-overhead aware. The dynamic VM consolidation problem is approached as a discrete combinatorial bi-objective problem of saving the most energy while keeping the migration cost at minimum. Factors affecting live migration cost and the parameters contributing to that cost are studied and an estimation model for migration overhead is proposed. Thereafter, a dynamic VM consolidation algorithm is designed and implemented based on a meta-heuristic algorithm, simulated annealing, that accounts for the migration cost imposed by a given consolidation plan. Simulation-based experiments are conducted on

CloudSim using real cloud workload traces from PlanetLab to evaluate the performance of the proposed algorithm. Results of the proposed algorithm are compared to those of Best Fit, First Fit and Worst Fit heuristic algorithms. For each experiment, the amount of energy consumed, number of consolidated servers and the number of carried out migrations are tracked. Among the compared algorithms, it was found that the migration-unaware first fit approach provides the least data center power consumption as well as number of released physical machines. The proposed migration overhead-aware simulated annealing based algorithm is found to consume almost the same amount of energy of that using a First Fit based consolidation. However, the proposed algorithm accounts for the cost due to live migration and is shown to reduce number of performed VM migrations by 10% compared to FF-based algorithm. ...

Contents

List of Tables

List of Figures

Abbreviations

ACO	**A**nt **C**olony **O**ptimization
ACS	**A**nt **C**olony **S**ystem
BF	**B**est **F**it
CPU	**C**entral **P**rocessing **U**nit
DVMC	**D**ynamic **V**irtual **M**achine **C**onsolidation
FF	**F**irst **F**it
FFD	**F**irst **F**it **D**ecreasing
HVMs	**H**osted **V**irtual **M**achines
IaaS	**I**nfrastructure **as a** **S**ervice
ICT	**I**nformation and **C**ommunication **T**echnology
IT	**I**nformation **T**echnology
MIPS	**M**illion **I**nstructions **P**er **S**econd
MLR	**M**ultiple **L**inear **R**egression
MMT	**M**inimum **M**igration **T**ime
MP	**M**igration **P**lan
NAS	**N**etwork **A**ttached **S**torage
OM	**O**verhead due to **M**igration
OS	**O**perating **S**ystem
PM	**P**hysical **M**achine

PaaS	**P**latform **as a S**ervice
PSO	**P**article **S**warm **O**ptimization
RC	**R**esource **C**apacity
RD	**R**esource **D**emand
RU	**R**esource **U**tilization
RC	**R**esource **C**apacity
RUA	**R**emaining **U**tilization **A**ware
SAN	**S**torage **A**rea **N**etwork
SaaS	**S**oftware **as a S**ervice
SLA	**S**ervice **L**evel **A**greement
VM	**V**irtual **M**achine
WF	**W**orst **F**it

Chapter 1

Introduction

1.1 Motivation

Cloud Computing has become a pervasive computing model that has grown from being an emerging trend to be "the norm" among many enterprises. Cloud data centers consume massive amounts of energy. According to a recent article about trends in data centers' energy consumption [1], the ICT (Information and Communication Technology) sector contributes to up to 2% of the carbon dioxide emissions worldwide. This number is almost equivalent to the emission amounts produced by the aviation industry [2]. Moreover, according to the same article, the IT (Information Technology) consumes around 7% of the total generated electricity globally. This number is envisioned to increase to 13% by 2030.

As stated by a report [3] published by Cisco [4] in 2017, Data centers account for 1.8% of all U.S. electrical consumption. Environmental organizations are pushing cloud providers to adopt environment-friendly policies to reduce their electricity consumption and reduce data centers carbon footprint. Energy efficient policies, also known as green policies, should be enforced to conserve energy, not only to increase the provider's profit margin, but also to protect the environment. In general, physical servers consume a lot of power even at low utilization levels. Data Centers are usually designed and configured to handle peak workloads. While in practice, those servers are rarely fully utilized, meaning that it is infrequent that all servers are needed to operate at their full operational power. Servers consume 80% of their peak power even at 20% utilization [5].

Server consolidation is one of the widely used approaches for energy efficiency in data centers. Live migration is used to relocate VMs from under-utilized servers to other servers without halting VM's operation. The objective is to free underutilized servers, turn them to a low power state mode and hence, save their operational power. However, most of the published works on server consolidation ignore the cost of live migration operation, which is in fact an expensive operation.

1.2 Objectives

The aim of this thesis is to investigate a dynamic virtual machine consolidation algorithm that is live migration overhead aware, based on the metaheuristic Simulated Annealing algorithm. This work tackles the estimation of the cost of live migration operation necessary to achieve consolidation and considering this estimated cost in dynamic virtual machine consolidation decisions. The objective of the proposed algorithm is to reduce the amount of power consumed by physical servers while keeping the overhead due to migration at minimum.

1.3 Methodology

In this thesis, the dynamic virtual machine consolidation problem is approached through modeling and simulating typical cloud environment setups using an existing discrete event simulation tool, Cloudsim [6]. The dynamic virtual machine consolidation is formulated as a discrete optimization problem. Simulated Annealing algorithm inspired by annealing in metallurgy, is utilized as a metaheuristic algorithm to solve the given optimization problem at hand. Moreover, a mathematical model for estimating the cost of live migration process is suggested and integrated with the consolidation algorithm. Extensive experiments are conducted and discussed, with varying simulation parameters and conditions.

1.4 Contribution

The key contributions of this thesis are as follows:

1. A mathematical formulation for the Dynamic Virtual Machine Consolidation
(DVMC) problem as a discrete optimization problem is proposed.

2. A model for estimating the overhead due to live migration based on suggested parameters is put forward.

3. A novel migration overhead-aware DVMC algorithm based on simulated annealing is proposed as a solution to the consolidation problem.

4. Experimental results based on extensive simulation across multiple algorithm parameters and conditions are conducted.

1.5 Thesis Organization

The rest of the thesis is organized as follows:

- Chapter 2 provides the background and discusses some technology-related concepts involved in this thesis.

- Chapter 3 presents a discussion for the related work as a comparative study

- Chapter 4 introduces the architecture, setup and used methodologies of the proposed algorithm.

- Chapter 5 presents the results and provides an analysis for the obtained results

- Chapter 6 draws a conclusion for the research work and discusses possible future research directions.