Evaluation of Energy-Aware Scheduling of Cloud-Computing Data-Center Tasks

Ziqian Dong
Department of Electrical
and Computer Engineering
New York Institute of Technology
ziqian.dong@nyit.edu

Roberto Rojas-Cessa
Department of Electrical
and Computer Engineering
New Jersey Institute of Technology
rojas@njit.edu

***Abstract*—In this position paper, we propose our plans to study energy efficiency and task scheduling schemes using the experimental facilities to be developed under the NSFCloud program. We outline the requirements of our experiments and the needed functions from the experimental facilities. To enable the study of fine granularity task scheduling, resource allocation, and performance monitoring, various virtualization technologies and flexible configuration to allow VM migration are desired functions from the experimental facilities. The access to experimental facilities will enable the study and validation of data center energy efficiency and resource management schemes.**

I. Introduction

Cloud computing is a growing paradigm for its high flexibility and scalability. The underlying structure supporting cloud computing is large data centers dedicated to housing a huge number of servers and providing them with infrastructure, software and computing resources. Energy consumed by these large-scale data centers contributes to the total cost of ownership of a data center as well as have impacts on environment. Our proposed research will focus on improving data center energy efficiency to ensure the sustainability and scalability of data centers.

We have proven theoretically and through simulation in our study that efficient task scheduling and finer granularity can improve data center efficiency and reduce energy consumption [1]–[4].

It is perceived that energy preservation is usually achieved at the cost of performance. However, from our previous research and simulated experiments, we found that energy efficient schedulers can be implemented to preserve energy with little impact on performance, and further tradeoff can still be made between energy and performance. A reconfigurable cloud experimental environment is a very desirable environment to produce more realistic results.

In our study, we used a simulator to compare our proposed Most Efficient Server First (MESF) scheduler with other schedulers to show the effectiveness of MESF [1], [3], [4]. Different schedulers are used to assign the same set of tasks to the same set of servers, and the produced schedule is used to calculate energy consumption of these schedulers.

II. The Challenge

One of the challenging issues for evaluation of different approaches for data center scheduling and energy analysis is the availability of real-world data for test and validation. We experimented with a trace dataset published by Google to study using virtual power consumption to evaluate task scheduling and resource allocation using real task requested and machine resource profile extracted from this dataset [4]. We would like to use a reconfigurable cloud experimental environment to verify our previously simulated result and continue to design and compare more energy aware schedulers.

III. The Requirements

The minimal requirements for us to deploy our experiment is to (1) have access to a set of servers, (2) define different types of workload based on distributions or real world trace data, (3) use a customized scheduler to assign the tasks in the generated workload to servers and (4) collect energy consumption data from servers. It would be ideal if dynamic power consumption of servers can be monitored and accessed, which can be used to produce accurate energy consumption measurements. However, if server energy profile and execution time of tasks are known, energy consumption of servers can be estimated as well. Besides energy consumption, we also wish to understand how performance of schedulers would be affected if their energy awareness is raised, so resource utilization monitoring is desired.

Further requirements may include virtualization of servers. Hypervisors such as KVM and Xen are useful tools for managing full machine virtualization [5] [6]. System level containers, such as LXC, are also alternatives to hypervisors [7]. We would use the Libvirt library for interacting with hypervisors and containers remotely [8].

VM migration is critical for data center scheduler and maintenance because it enables task consolidation, which in turn gives chances to turn off servers. However, VM migration is a complicated issue that involves VM checkpointing, data transport and VM restoration. It has an overhead that can not be ignored. In order to fully understand efficiency of VM migration, delays on all three of these stages have to be monitored. We would use the Chameleon cloud to study VM migration, the overhead reduction and timing efficiency.

REFERENCES

[1] N. Liu, Z. Dong, and R. Rojas-Cessa, "Task and server assignment for reduction of energy consumption in datacenters," in *the 11th IEEE International Symposium on Network Computing and Applications, IEEE NCA12*, August 2012, pp. 1–4.

[2] R. Rojas-Cessa, S. Pessima, and T. Tian, "Experimental evaluation of energy savings of virtual machines in the implementation of cloud computing," in *Proc. IEEE WOCC 2012*, April 2012, pp. 1–6.

[3] N. Liu, Z. Dong, and R. Rojas-Cessa, "Task scheduling and server provisioning for energy-efficient cloud-computing data centers," in *the 3rd International Workshop on Data Center Performance, DCPerf 2013*, July 2013, pp. 1–4.

[4] Z. Dong, W. Zhuang, and R. Rojas-Cessa, "Energy-aware scheduling schemes for cloud data centers on google trace data," in *the 2014 IEEE Online Conference on Green Communications*, November 2014, pp. 1–6.

[5] "The Xen Project, the powerful open source industry standard for virtualization," http://www.xenproject.org/, accessed: 2014-11-1.

[6] "KVM," www.linux-kvm.org/page/Main_Page, accessed: 2014-11-1.

[7] "LCX - Linux Containers," https://linuxcontainers.org/, accessed: 2014-11-1.

[8] "libvirt: The virtualization api," http://libvirt.org/index.html, accessed: 2014-11-1.