

Energy Efficient Allocation of Virtual Machines in Cloud Data Centers

Anton Beloglazov and Rajkumar Buyya
Cloud Computing and Distributed Systems (CLOUDS) Laboratory
Department of Computer Science and Software Engineering
The University of Melbourne, Australia
 {abe, raj}@csse.unimelb.edu.au

Abstract—Rapid growth of the demand for computational power has led to the creation of large-scale data centers. They consume enormous amounts of electrical power resulting in high operational costs and carbon dioxide emissions. Moreover, modern Cloud computing environments have to provide high Quality of Service (QoS) for their customers resulting in the necessity to deal with power-performance trade-off. We propose an efficient resource management policy for virtualized Cloud data centers. The objective is to continuously consolidate VMs leveraging live migration and switch off idle nodes to minimize power consumption, while providing required Quality of Service. We present evaluation results showing that dynamic reallocation of VMs brings substantial energy savings, thus justifying further development of the proposed policy.

Keywords—Energy efficiency; Cloud computing; Energy consumption; Green IT; Resource management; Virtualization; Allocation of virtual machines; Live migration of virtual machines.

I. INTRODUCTION

In recent years, IT infrastructures continue to grow rapidly driven by the demand for computational power created by modern compute-intensive business and scientific applications. However, a large-scale computing infrastructure consumes enormous amounts of electrical power leading to operational costs that exceed the cost of the infrastructure in few years. For example, in 2006 the cost of electricity consumed by IT infrastructures in US was estimated as 4.5 billion dollars and tends to double by 2011 [1]. Except for overwhelming operational costs, high power consumption results in reduced system reliability and devices lifetime due to overheating. Another problem is significant CO_2 emissions that contribute to the greenhouse effect.

One of the ways to reduce power consumption by a data center is to apply virtualization technology. This technology allows one to consolidate several servers to one physical node as Virtual Machines (VMs) reducing the amount of the hardware in use. Recently emerged Cloud computing paradigm leverages virtualization and provides on-demand resource provisioning over the Internet on a pay-as-you-go basis [2]. This allows enterprises to drop the costs of maintenance of their own computing environment and outsource the computational needs to the Cloud. It is essential for Cloud providers to offer reliable Quality of Service (QoS) for the customers that is negotiated in terms of

Service Level Agreements (SLA), e.g. throughput, response time. Therefore, to ensure efficient resource management and provide higher utilization of resources, Cloud providers (e.g. Amazon EC2) have to deal with power-performance trade-off, as aggressive consolidation of VMs can lead to performance loss.

In this work we leverage live migration of VMs and propose heuristics for dynamic reallocation of VMs according to current resources requirements, while ensuring reliable QoS. The objective of the reallocation is to minimize the number of physical nodes serving current workload, whereas idle nodes are switched off in order to decrease power consumption. A lot of research has been done in power efficient resource management in data centers, e.g. [3], [4]. In contrast to previous studies, the proposed approach can effectively handle strict QoS requirements, heterogeneous infrastructure and heterogeneous VMs. The algorithms are implemented as fast heuristics; they do not depend on a particular type of workload and do not require any knowledge about applications executing on VMs.

II. REALLOCATION HEURISTICS

Allocation of VMs can be divided in two: the first part is admission of new requests for VM provisioning and placement VMs on hosts, whereas the second part is optimization of current allocation of VMs. The first part can be considered as a bin packing problem with variable bin sizes and prices. To solve it we apply modification of the Best Fit Decreasing (BFD) algorithm. In our modification (MBFD) we sort all VMs in decreasing order of current utilization and allocate each VM to a host that provides the least increase of power consumption due to this allocation. This allows to leverage heterogeneity of the nodes by choosing the most power-efficient ones. The complexity of the allocation part of the algorithm is $n \cdot m$, where n is the number of VMs that have to be allocated and m is the number of hosts.

Optimization of current allocation of VMs is carried out in two steps: at the first step we select VMs that need to be migrated, at the second step chosen VMs are placed on the host using MBFD algorithm. We propose four heuristics for choosing VMs to migrate. The first heuristic, Single Threshold (ST), is based on the idea of setting upper utilization threshold for hosts and placing VMs while keeping

the total utilization of CPU below this threshold. The aim is to preserve free resources in order to prevent SLA violation due to consolidation in cases when resource requirements by VMs increase. At each time frame all VMs are reallocated using MBFD algorithm with additional condition of keeping the upper utilization threshold not violated. New placement is achieved by live migration of VMs.

The other three heuristics are based on the idea of setting upper and lower utilization thresholds for hosts and keeping total utilization of CPU by all VMs between these thresholds. If the utilization of CPU for a host falls below the lower threshold, all VMs have to be migrated from this host and the host has to be switched off in order to eliminate idle power consumption. If the utilization goes over the upper threshold, some VMs have to be migrated from the host to reduce utilization to prevent potential SLA violation. We propose three policies for choosing VMs that have to be migrated from the host: (1) Minimization of Migrations (MM) – migrating the least number of VMs to minimize migration overhead; (2) Highest Potential Growth (HPG) – migrating VMs that have the lowest usage of CPU relatively to requested in order to minimize total potential increase of the utilization and SLA violation; (3) Random Choice (RC) – migrating the necessary number of VMs by picking them according to a uniformly distributed random variable.

III. EVALUATION

The proposed heuristics have been evaluated by simulation using CloudSim toolkit [5]. The simulated data center comprises 100 heterogeneous physical nodes. Each node is modeled to have one CPU core with performance equivalent to 1000, 2000 or 3000 MIPS, 8 GB of RAM and 1 TB of storage. Users submit requests for provisioning of 290 heterogeneous VMs that fill the full capacity of the data center. For the benchmark policies we simulated a Non Power Aware policy (NPA) and DVFS that adjusts the voltage and frequency of CPU according to current utilization. We present results obtained using ST policy and the best two-threshold policy, MM policy. Besides that, the policies have been evaluated with different values of the thresholds.

The simulation results presented in Table I show that dynamic reallocation of VMs according to current utilization of CPU brings higher energy savings compared with static allocation policies. MM policy achieves the best energy savings: by 83%, 66% and 23% less energy consumption relatively to NPA, DVFS and ST policies respectively with thresholds 30-70% and ensuring percentage of SLA violations of 1.1%; and by 87%, 74% and 43% with thresholds 50-90% and 6.7% of SLA violations. MM policy leads to more than 10 times fewer VM migrations than ST. The results show the flexibility of the proposed algorithms, as the thresholds can be adjusted according to SLA requirements. Strict SLA (1.11%) allow achievement of the energy consumption of

1.48 KWh. However, if SLA are relaxed (6.69%), the energy consumption is further reduced to 1.14 KWh.

Table I
SIMULATION RESULTS

Policy	Energy	SLA	Migr.	Avg. SLA
NPA	9.15 KWh	-	-	-
DVFS	4.40 KWh	-	-	-
ST 50%	2.03 KWh	5.41%	35 226	81%
ST 60%	1.50 KWh	9.04%	34 231	89%
MM 30-70%	1.48 KWh	1.11%	3 359	56%
MM 40-80%	1.27 KWh	2.75%	3 241	65%
MM 50-90%	1.14 KWh	6.69%	3 120	76%

IV. CONCLUSION

In this work we have proposed and evaluated heuristics for dynamic reallocation of VMs to minimize energy consumption, while providing reliable QoS. The obtained results show that the technique of dynamic reallocation of VMs and switching off the idle servers brings substantial energy savings and is applicable to real-world Cloud data centers. For the future work, we propose to investigate the consideration of multiple system resource in reallocation decisions, such as network interface and disk storage, as these resources also significantly contribute to the overall energy consumption. Other interesting directions for the future work are investigation of setting the utilization thresholds dynamically according to a current set of VMs allocated to a host, leveraging multi-core CPU architectures, and decentralization of the optimization algorithms to improve scalability and fault tolerance. Besides the reduction of operational and establishment costs, the work has social significance as it decreases carbon dioxide footprints and energy consumption by modern IT infrastructures.

REFERENCES

- [1] R. Brown *et al.*, "Report to congress on server and data center energy efficiency: Public law 109-431," *Lawrence Berkeley National Laboratory*, 2008.
- [2] R. Buyya, C. S. Yeo, and S. Venugopal, "Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities," in *Proceedings of HPCC'08*. IEEE CS Press, Los Alamitos, CA, USA, 2008.
- [3] D. Kusic, J. O. Kephart, J. E. Hanson, N. Kandasamy, and G. Jiang, "Power and performance management of virtualized computing environments via lookahead control," *Cluster Computing*, vol. 12, no. 1, pp. 1–15, 2009.
- [4] S. Srikantaiah, A. Kansal, and F. Zhao, "Energy aware consolidation for cloud computing," *Cluster Computing*, vol. 12, pp. 1–15, 2009.
- [5] R. Buyya, R. Ranjan, and R. N. Calheiros, "Modeling and simulation of scalable cloud computing environments and the CloudSim toolkit: Challenges and opportunities," in *Proceedings of HPCS'09*. IEEE Press, NY, USA, 2009.