

Power-Aware Virtual Machine Scheduling-policy for Virtualized Heterogeneous Multicore Systems

Taranpreet Kaur, Inderveer Chana

Abstract— This paper presents a systematic approach to correctly provision server resources in data centers, resulting in minimum energy consumption and SLA violations. In particular, we describe a hybrid method for efficient server provisioning in virtualized heterogeneous multicore Cloud data centers. The objective is to place VMs on host while keeping total utilization of CPU below defined threshold and then optimizing the VM allocation by constrained consolidation of VMs leveraging live migration and switching off idle nodes to minimize power consumption. As excessive migration of VMs can lead to significant SLA violation which can hamper required Quality of Service, in proposed approach not only overall number of used servers are minimized, but also number of migrations are minimized, resulting in minimum energy consumption and least SLA violations. We present evaluation results showing that our approach can provide up to 13% of energy savings, reduced SLA violations and fewer VM migrations as compared to existing techniques, while preventing frequent power cycling of servers, thus justifying further development of stated policy.

Keywords— Green IT; Cloud computing; Energy efficiency; Allocation of virtual machines; Resource management; Virtualization; Dynamic consolidation

I. INTRODUCTION

Data centers are very expensive to operate due to the power and cooling requirements of IT equipment like network switches, storage and servers. As demand for IT services increases, the power required to operate data centers also increases. The EPA estimated that energy consumption in data centers exceeded 100 billion kWh in 2011, at a cost of \$7.4 billion [1]. Furthermore, it is estimated that servers consume 0.5 percent of the world's total electricity usage [2], which if current demand continues, is projected to quadruple by 2020. Rising energy costs, regulatory requirements and social concerns over green house gas emissions have made reducing power consumption critical to data center operators. However, energy efficiency is for nothing if the data center cannot deliver IT services according to predefined QoS goals, because SLA violations can result in lost business revenue.

Taranpreet Kaur
Cse Dept.,Thapar University,
Patiala, India

Inderveer Chana
Cse Dept.,Thapar University,
Patiala, India

For example, Amazon observed that every additional 100 ms of latency costs them a 1% loss in sales, similarly Google found that an extra 500 ms in search page generation time reduced traffic by nearly 20% [3]. Hence, another challenge data center operators face is provisioning of IT resources such that SLA violations are minimized and business revenue loss is prevented. Today, SLA violations are often avoided by over-provisioning IT resources but this leads to excessive energy consumption, and may also cause increased expenditure due to capital overhead, maintenance costs, etc. Thus, presently an important question in data center resource management is how to correctly provision IT equipment, so that SLA violations are avoided and energy consumption is reduced.

Correct provisioning of resources is a difficult task because of variations in workload demands. Most data center workload demands are very bursty in nature and often vary considerably even during the course of a single day. This makes it difficult to provision resources efficiently. A single size i.e. static provisioning cannot fit all and can result in either over-provisioning or under provisioning of resources.

Berl *et al.* [4] provide four approaches for increasing energy efficiency i.e. (i) hardware level energy optimization, (ii) energy-aware scheduling in grid systems, (iii) server consolidation by means of virtualization and (iv) power minimization in wireless and wired networks. The solution proposed in this paper is based on third approach that deals with consolidation of applications to fewer servers up to some constrained level on virtualized cluster data centers taking in account, heterogeneity in workload demands and multiple system resources while reallocation decisions. Lot of research has been conducted in power efficient resource management in data centers, e.g. [5], [6], [7].

In contrast to previous studies, the proposed virtual machine provisioning approach can effectively handle strict QoS requirements, heterogeneous VMs and heterogeneous infrastructure. The algorithms implemented do not depend on a particular type of workload and do not require any knowledge about applications executing on VMs. Our main contributions are:

1. We present a detailed characterization study of existing VM allocation and optimization techniques considering variations in workload demands.
2. We present evaluation results and discovered level of energy consumption and SLA violation in existing techniques.

3. We propose a hybrid strategy to provision IT resources: A systematic provisioning approach to first place VMs on hosts while keeping total utilization of CPU below defined threshold and then optimizing the VM allocation by continuously consolidating VMs leveraging live migration according to current requirements. A coordinated management of this hybrid approach achieves a significant improvement in energy efficiency while ensuring reliable QoS.
4. We evaluate our server provisioning system considering heterogeneous workload and heterogeneous multiple system resources. The results reveal that our hybrid provisioning system is superior to existing approaches in terms of meeting the system's SLA requirements while conserving power and reducing provisioning cost.

Rest of the paper is organized as follows. Section II presents analysis of available energy efficient allocation and optimization heuristics for virtual machines in Cloud data centers. Section III describes the proposed hybrid provisioning approach for energy efficient virtual machine scheduling in heterogeneous infrastructure. Section IV discusses experimental evaluation of proposed as well as existing VM scheduling approaches and finally Section V concludes the paper with a summary of our research work and a description of future directions.

II. REALLOCATION HEURISTICS ANALYSIS

Virtual machine allocation can be divided in two: first is admission of new requests for VM provisioning and then placement of VMs on available hosts, while the second part is optimization of current VM allocation on hosts. In the existing VM allocation policies [8], admission of new requests for VM provisioning is considered as a bin packing problem with varying bin sizes and prices. The bin-packing problem is a type of NP-hard problem, and number of heuristic algorithms are available to give sub-optimal results, such as Next-Fit, First-Fit, Best-Fit and Best-Fit Decreasing that is shown to use no more than $11/9 \cdot \text{OPT} + 1$ bins (where OPT is the number of bins given by the optimal solution) [9]. Modified Best Fit Decreasing (MBFD) algorithm [5] allows leveraging heterogeneity of the nodes by first sorting all VMs in decreasing order of current utilization and allocating each virtual machine to a host that provides least increase of power consumption due to this allocation policy. Complexity of this modified BFD allocation part is $m \cdot n$, where m is the number of VMs to be allocated and n is the number of hosts available.

Optimization of current allocation of VMs on hosts can be carried out by dynamic migration of VMs according to the performance requirements. If VMs do not utilize all the provided resources, then they can be logically resized and consolidated to the minimum number of physical nodes.

Resulting idle nodes can be switched to sleep mode which helps in elimination of idle power consumption and reduction in the total energy consumption by the data center. Existing optimization techniques are based on following two criteria:

1. Setting upper utilization threshold for hosts and keeping total utilization of CPU by all VMs below this threshold.
2. Setting upper and lower utilization thresholds for hosts and keeping total utilization of CPU between these thresholds.

If the utilization of CPU for a host falls below the lower utilization threshold then all VMs have to be migrated from this under-utilized host and the host has to be switched to sleep mode in order to eliminate idle power consumption and if the utilization goes over the upper threshold, then some VMs have to be migrated from the over-utilized host to reduce utilization and to prevent potential SLA violation. Existing optimization approaches are discussed below:

1. Single threshold (ST): It is based on the idea of setting upper utilization threshold for hosts and placing VMs on hosts keeping the total utilization of CPU below the stated threshold. The aim is to preserve free resources in order to avoid SLA violation due to consolidation in those cases when resource requirements by VMs increase. New placement of VMs is achieved by the live migration of VMs.
2. Minimization of Migrations (MM): This policy is based on criteria of selection of minimum number of VMs required to migrate from a host to lower the CPU utilization by the host below the upper utilization threshold in case the upper threshold is violated.
3. Highest Potential Growth (HPG): If the upper utilization threshold is violated, this policy migrates VMs that have the lowest usage of the CPU comparatively to the CPU capacity defined by the VM parameters. This helps in minimization of potential increase in the host's utilization and prevention of SLA violation.
4. Random Choice (RC): This policy focuses on random selection of number of VMs required to decrease CPU utilization by a host below the stated upper utilization threshold for host.

Presently, resource allocation in a Cloud data center intends to provide high performance while meeting SLAs, without focusing on allocation of VMs to minimize energy consumption. To investigate both performance and energy efficiency, three critical issues must be addressed which are discussed as follows:

1. Excessive power cycling of a server can cause reduction in its reliability.

2. Turning resources off in a dynamic environment can be risky from the QoS perspective. Due to the inconsistency in the workload and aggressive consolidation of VMs, some VMs may not obtain required resources under peak load situation, and fail to meet the desired QoS.
3. Ensuring SLAs brings challenges to accurately manage application performance in virtualized environments.

All these issues require effective consolidation policies which can minimize energy consumption without negotiating the user-specified QoS requirements. In next section, we are discussing our hybrid approach for VM allocation to deal with the issues mentioned above.

III. HYBRID PROVISIONING APPROACH

The intuition behind our approach is that first place VMs on host while keeping total utilization of CPU below defined threshold and then optimizing VM allocation by constrained migration of VMs with host utilization below the lower utilization threshold, leveraging live migration and switching off idle nodes to minimize power consumption. As mentioned in above section that excessive migration of VMs can lead to significant SLA violation hampering the required Quality of Service, the proposed approach can assist in correctly provisioning server resources in data center by optimal VM allocation and constrained migration of VMs, resulting in least SLA violations and minimum energy consumption.




Fig. 1. Hybrid provisioning architecture

Fig. 1 depicts the conceptual architecture for the hybrid provisioning approach. Our solution comprises of two components. First component helps in preserving free resources by placing VMs below defined static threshold in order to prevent SLA violation due to consolidation in cases when resource requirements by VMs increase and second component helps in reducing level of power consumption by migrating VMs from only those hosts having total CPU utilization below lower migration threshold and switching off the spare host to eliminate idle power consumption.

The pseudo-code for components of the hybrid algorithm is presented in Algorithm 1 and 2. The algorithm 1 sets an upper utilization threshold for hosts. Then, for each VM in the list of VMs, it repeatedly looks through the list of hosts and finds a host with total CPU utilization below the static utilization threshold. Then chosen host is set as the allocated host and VM is placed on it. The increase in host power after allocation of VMs is added to the total power consumption. This procedure is repeated for each VM in the VM list and algorithm stops when all of the VMs are placed on available hosts. The complexity of the algorithm is proportional to the product of the number of VMs that have to be allocated and the number of hosts available.

Optimization of current allocation of VMs on hosts is done using Algorithm 2. In algorithm 2, first the list of hosts is sorted in ascending order according to total utilization of CPU by all VMs. Then lower utilization threshold also known as migration threshold is set to constrain migration of VMs. A counter equal to total number of hosts in sorted host list is set to start the optimization procedure. First, for each host i in sorted host list, if host utilization lies below migration threshold i.e. host is under-utilized, then VMs on that particular host are located for migration. Migration of VMs is first checked at host j , host with highest CPU utilization i.e. last host in sorted host list. Migration to the host j is done if total CPU utilization of host after VM consolidation remains below static upper threshold defined in Algorithm 1. If total CPU utilization of host j will go across upper utilization threshold when VM consolidation will be done, then migration of VMs from under-utilized host to this particular host is not possible, therefore preceding host in sorted host list is checked for migration and counter is decreased by one. This procedure is carried out till all the VMs on host i are migrated to host j . This procedure repeats for all under-utilized hosts till host value i become equal to the value of counter i.e. when there is no host left to which migration of VMs of under-utilized hosts can be done. In this way, this approach constrains the migration of VMs, such that SLA violations are minimized and required Quality of service is attained.

Algorithm 1: Initial Allocation Procedure

```

1 Input: hostList, vmList Output: allocation of VMs
2 hostList.setUpperThresholdUtilization ()
3 allocatedHost <- NULL
4 foreach vm in vmList do

```

```

5   foreach host in hostList do
6     if hostUtil < Thresh_Upper then
7       allocatedHost <- host
8       TotalPower <-HostPower
9     if allocatedHost != NULL then
10      allocate vm to allocatedHost
11 return allocation
12 optimize vm allocation with Algorithm 2

```

Algorithm 2: Optimization procedure

```

1 Input: vmAllocationList, hostList Output: migrationList
2 sort hostList according to CPU utilization in ascending order
  and set Migration threshold
3 i=1, j=N
4 foreach  $h_i$  in sortedHostlist do
5   while  $h_i$ .util_rate < Migration_Thresh and  $i < j$  do
6     if  $h_i$ .util_rate +  $h_j$ .util_rate <  $h_j$ .Thresh_Upper then
7       migrate  $vm_i$  to  $h_j$ 
8       sparehostList.add( $h_i$ )
9       usedhostList.add( $h_j$ )
10      break
11     else
12        $j = j - 1$ ;
13 return migrationList

```

In next section, the proposed hybrid VM provisioning approach is evaluated in comparison with existing allocation policies.

IV. EXPERIMENTAL EVALUATIONS

The proposed approach has been evaluated by simulation using CloudSim toolkit [10]. The simulated data center comprises 200 heterogeneous physical nodes. Each node is modelled to have dual CPU core with performance equivalent to 1860 or 2660 MIPS, 5 GB of RAM and 1 TB of storage. Users submit request for provisioning of 310 heterogeneous VMs that fills the full capacity of the data center. For the benchmark policies, we also stimulated DVFS that adjusts the frequency and voltage of CPU according to current utilization and a Non Power Aware policy (NPA). Table I presents results obtained from single threshold policy i.e. ST policy, the best two-threshold policy i.e. MM policy and benchmark policies. Results obtained using the hybrid approach is presented in Table II. Moreover, policies have been evaluated with different threshold values.

TABLE I. Simulation results for existing policies

Policy	Energy	SLA Violation	Total VM Migration	Avg. SLA Violation
NPA	150.68 KWh	-	-	-
DVFS	71.40 KWh	-	-	-
Static Threshold (ST) 80%	52.98 KWh	4.49%	20231	13.65%
Static Threshold (ST) 90%	52.75 KWh	8.41%	19676	27.04%
Minimum Migration (MM) 40-80%	49.67 KWh	3.11%	9800	12.89%
Minimum Migration (MM) 50-90%	49.25 KWh	5.25%	9183	16.35%

TABLE II. Simulation results for hybrid provisioning approach

Upper Utilization threshold	Migration threshold	Energy	SLA Violation	Total VM Migration	Avg. SLA Violation
90%	20%	50.00 KWh	4.30%	9900	13.33%
90%	30%	48.81 KWh	2.72%	8883	12.16%
90%	40%	47.19 KWh	2.29%	7474	12.04%
90%	50%	42.71 KWh	3.01%	6351	12.65%

Comparing the simulation results presented in Table I and II, it can be concluded that hybrid approach for VM provisioning provides better energy savings:

- i. by 67%, 31%, 7% and 0.89% less energy consumption relatively to NPA, DVFS, ST and MM policies with upper utilization threshold 90% and lower utilization threshold 30% and ensuring percentage of SLA violations of 2.72%; and
- ii. by 68%, 34%, 10% and 4% less energy consumption in compare to NPA, DVFS, ST and MM policies with thresholds 40-90% and 2.29% of SLA violations; and
- iii. by 71%, 40%, 19% and 13% less energy consumption comparatively to NPA, DVFS, ST and MM policies with thresholds 50-90% and 3.01% of SLA violations.

From Table I and II, it is evident that hybrid provisioning approach leads to more than 3 times fewer VM migrations than ST policy and more than one time fewer VM migrations than MM policy. The results prove the flexibility of the proposed algorithm, as the utilization thresholds can be adjusted according to SLA requirements. Strict SLA (2.29%) allow achievement of the energy consumption of 47.19KWh but if SLA are relaxed (3.01%), then energy consumption is further reduced to 42.71 KWh.

V. CONCLUDING REMARKS AND FUTURE DIRECTIONS

This work precedes the Cloud computing field in two ways. First, it has a vital role in the reduction of energy consumption costs in Cloud data centers, and thus helps in developing a strong and competitive Cloud computing industry. Second, consumers are gradually becoming more conscious about the system environment. A recent study shows that data centers represent a huge and rapidly growing energy consumption division of the economy and a significant source of CO2 emissions [11]. Key energy policy of many countries around the world focuses on reduction of greenhouse gas emissions. We have presented and evaluated a hybrid energy-aware resource allocation approach making use of dynamic consolidation of VMs. The experiment results have revealed that this approach results in considerable reduction of energy consumption in Cloud data center relatively to former techniques. For future work, we propose to investigate setting of utilization thresholds dynamically according to current allocation of VMs to a host, influencing multi-core architectures, and decentralization of the optimization algorithms for improving scalability and fault tolerance of Cloud data centers. Further interesting directions for the future work are investigation of effects of network interface and disk storage on reallocation decisions, as both significantly contribute to the overall energy consumption.

ACKNOWLEDGMENT

The research leading to these results was supported by AICTE under Research Promotion Scheme Ref. No.8023/RID/RPS-151 (pvt)/2011-2012.

REFERENCES

- [1] Report to congress on server and data center energy efficiency, Public law, U.S. Environmental Protection Agency, ENERGY STAR Program, 2007.
- [2] W. Forrest, "How to cut data centre carbon emissions?," <http://www.computerweekly.com/Articles/2008/12/05/233748/how-to-cut-data-centre-carbon-emissions.htm>, December 2008.
- [3] G. Linden, "Make data useful," <http://www.scribd.com/doc/4970486/Make-Data-Useful-by-Greg-Linden-Amazoncom>, 2006.

- [4] A. Berl, et al., "Energy-Efficient Cloud Computing," Elsevier's The Computer Journal, 2009.
- [5] D. Kusic, et al., "Power and performance management of virtualized computing environments via lookahead control," Elsevier Journal of Cluster Computing, vol. 12, no. 1, pp. 1-15, 2009.
- [6] S. Srikantaiah, A. Kansal, F. Zhao, "Energy aware consolidation for cloud computing," Elsevier Journal of Cluster Computing, vol. 12, pp. 1-15, 2009.
- [7] A. Beloglazov, R. Buyya, "Energy efficient allocation of virtual machines in cloud data centers," Proc. 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing, pp. 577-578. IEEE Press, 2010.
- [8] A. Beloglazov, J. Abawajy, R. Buyya, "Energy-aware resource allocation heuristics for efficient management of data centers for Cloud computing," Elsevier Journal of Future Generation Computer Systems, vol. 28, pp. 755-768, 2012.
- [9] M. Yue, "A simple proof of the inequality $FFD(L) < 11/9 OPT(L) + 1$, for all l for the FFD bin-packing algorithm," Acta Mathematicae Applicatae Sinica (English Series), vol. 7, no. 4, pp. 321-331, 1991.
- [10] R. Buyya, R. Ranjan, R. N. Calheiros, "Modelling and simulation of scalable cloud computing environments and the CloudSim toolkit: Challenges and opportunities," Proceedings of HPCS'09. IEEE Press, NY, USA, 2009.
- [11] P. Johnson, T. Marker, "Data centre energy efficiency product profile," Tech. Rep., 2009/05, 2009.

International and National Conferences. She is the Member of Centre of Excellence in Grid Computing at Thapar University and Coordinator of Infosys for Campus Connect Foundation Program.



Taranpreet kaur received Btech degree in Computer Science and Engineering from Punjab Technical University in 2011. She is currently pursuing M.E. in Software Engineering from Thapar University, Patiala (India). Her research interests

focus on energy efficient resource management and scheduling in distributed systems, especially in Grids and Clouds.



Inderveer Chaa is Ph.D in Computer Science with specialization in Grid Computing and M.E. in Software Engineering from Thapar University and B.E. in Computer Science and Engineering. She joined Thapar University in 1997 as Lecturer and has over

thirteen years of experience. She is presently working as Associate Professor in Computer Science and Engineering Department of Thapar University. Her research interests include Grid computing and Cloud Computing and other areas of interest are Software Engineering and Software Project Management. She has more than 50 research publications in reputed journals and conferences. She has attended numerous