REDUCING POWER CONSUMPTION & DELAY AWARE RESOURCE ALLOCATION IN CLOUD DATACENTERS

[1] MR. A. A. PATEL, [2]PROF. J. N. RATHOD

[1]M.E.[Computer] Student, Department Of Computer Engineering, Atmiya Institute of Technology and Science, Rajkot, Gujarat
[2]Professor Department of Computer and Information Technology, Atmiya Institute of Technology and Science, Rajkot, Gujarat

patelambrish9@gmail.com,jnrathod@aits.edu.in

***ABSTRACT*:** *Continuous growth and demand in computational power endorsed for the creation of high scale data centers, requires very large amount of electrical power which acquires high operational cost and emission of carbon dioxide. Virtualization is used to reduce power consumption by providing number of virtual servers onto a smaller amount of computing resources. Current Cloud Computing scenario focuses on high Quality of Service for consumers which results in deal with power performance factor. In this paper, an efficient resource management policy for Virtualized Cloud data centers has been proposed. Reduction of power consumption can be done by switching off idle nodes and leveraging live migration of VMs. In addition, operational cost can be affected by dynamic reallocation of VMs and network propagation delay. Aim of this paper is to optimize the power consumption of VMs and to reduce network delay.*

Keywords—Cloud; Virtualization; Power Consumption;

I: INTRODUCTION

1.Cloud: "Cloud Computing is defined by a large-scale distributed computing paradigm that is driven by economies of scale, in which a pool of abstracted, virtualized, dynamically-scalable, managed computing power, storage, platforms, and services are delivered on demand to external customers over the Internet".

In a simple, topological sense, a cloud computing solution is made up of several elements: clients, the datacenters and distributed servers.

2.Virtulaization: In Virtualization, more than one VM can run on a single host executing entirely different applications. So whenever a request arrives for an application, a new virtual machine is created in the data center and is placed on one of the physical host machine to run the requested application.

Reducing the energy consumption of data centers is a challenging and complex task because computing applications and data are growing so quickly that even larger servers and disks are needed to process them fast enough so that it meets the time constraints provided by users.

II: EXISTING POWER CONSUMPTION IN CLOUD COMPUTING

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 powerefficient ones. The complexity of the allocation part of the algorithm is n_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 utilizationthreshold for hosts and placing VMs while keeping the total utilization of CPU below this threshold. The aim isto preserve free resources in order to prevent SLA violationdue to consolidation in cases when resource requirements by VMs increase. At each time frame all VMs are reallocatedusing MBFD algorithm with additional condition of keepingthe upper utilization threshold not violated. New placementis 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. The other heuristic 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 themaccording to a uniformly distributed random variable. The simulation table of this heuristic is as below.
The proposed heuristics have been evaluated by simulation using CloudSim toolkit [5]. The simulated datacenter 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 energyconsumption is further reduced to 1.14 KWh.

Policy	Energy	SLA	Migr.	Av.SLA
NPA	9.15kwh	-	-	-
DVFS	4.40kwh	-	-	-
ST 50%	2.03kwh	5.41%	35 226	81%
ST 60%	1.50kwh	9.04%	34 231	89%
MM30-70%	1.48kwh	1.11%	3 369	56%
MM 40-80%	1.27kwh	2.75%	3 241	65%
MM 50-90%	1.14kwh	6.69%	3 120	76%

III: PROPOSED SYSTEM
We can state the problem as follow:

"Derive an algorithmic solution for an energy-efficient resource management system for virtualized Cloud data centers that reduces operational costs by reducing power consumption and delay while provides required Quality of Service (QoS)."

In this dissertation, I have made an attempt to design and implement framework for live migration and dynamic reallocation of VMs according to current utilization, while ensuring reliable QoS and minimize power consumption and delay using the following routines:

1. Modified Best Fit Decreasing Algorithm
2. Genetic Algorithm

Sample Migration algorithms:

- Input hostList Output migrationList
- For each host in hostList do
- vmList <- host.getvmList()
- vmList.sortIncreasingUtilization()
- hUtil<-host.getUtil()
- While hUtil > Thresh_up do
- bestFitvm <- NULL
- for each vm in vmList do
- bestFitvm <- vm
- if vm.getUtil() > hUtil-Thresh_up then
- break
- if bestFritvm != NULL then
- hUtil <- hUtil –bestFitVm.Util
- migrationList.add(bestFitVm.Util
- VmList.remove(bestFitVm)
- if hUtil < Thresh_low then
- migrationList.add(host.getVmList())
- VmList.remove(host.getVmlist())
- do Host.shutdown <- TRUE
- return migrationList

The pseudo-code for the MM algorithm for the over-utilization case is presented in Algorithm. The algorithm sorts the list of VMs in the increasing order of the CPU utilization. Then, it repeatedly looks through the list of VMs and finds a VM that is the best to migrate from the host. The best VM is the one removal of which leaves the total utilization of host by all VMs just below the upper utilization threshold with minimum of margin. If there is no such a VM, the algorithm selects the last VM of the list which is of the highest utilization, removes it from the list of VMs, and proceeds to a new iteration. The algorithm stops when the new utilization of the host is below the upper utilization threshold. The complexity of the algorithm is n * m, where n is the number of over-utilized hosts and m is the number of VMs allocated to these hosts. If the lower threshold is violated then all the VMs from the host is migrated and host is shutdown to reduce idle power consumption.

VI: CONCLUSION

"Large scale data-centers which consume enormous amount of electrical power resulting in high operational cost and emission of carbon dioxide".

Virtualization is a promising approach to reduce this power consumption by consolidating multiple virtual servers onto a smaller number of computing resources. But modern Cloud computing environments have to provide high Quality of Service for their customers resulting in the necessity to deal with power-performance trade-off.

REFRENCES

[1] R. Nathuji and K. Schwan, "VirtualPower: Coordinated power management in virtualized enterprise systems," *ACM SIGOPS Operating Systems Review 2007*; volume 41 issue 6, December 2007, pp. 265–278.

[2] D. Kusic, J. Kephart, J. Hanson, N. Kandasamy, G. Jiang, "Power and performance management of virtualized computing environments via lookahead control," *Cluster Computing 2009*,volume 12 number 1, pp. 1–15.

[3] S. Srikantaiah, A. Kansal, F. Zhao, "Energy aware consolidation for cloud computing," in *Proceedings of the 2008 conference on Power aware computing and systems*, USENIX Association Berkeley, CA, USA, 8-10 December 2008, pp. 10.

[4] M. Cardosa, M. Korupolu, A. Singh, "Shares and utilities based power consolidation in virtualized server environments," in *Proceedings of the 11th IFIP/IEEE Integrated Network Management (IM 2009)*, Long Island, NY, USA, 1-5 June 2009, pp. 327-334.

[5] A. Verma, P. Ahuja, A. Neogi, "pMapper: Power and migration cost aware application placement in virtualized systems," in *Proceedings of the 9th ACM/IFIP/USENIX International Conference on Middleware (Middleware 2008)*, Springer, Leuven, Belgium, 1-5 December 2008, pp. 243–264.

[6] A. Verma, G. Dasgupta, T. Nayak, P. De, R. Kothari, "Server workload analysis for power minimization using consolidation," in *Proceedings of the 2009 USENIX Annual Technical Conference*, San Diego, CA, USA, 14-19 June 2009, pp. 28.,.

[7] A. Gandhi, M. Harchol-Balter, R. Das, C. Lefurgy, "Optimal power allocation in server farms," in *Proceedings of the 11th International Joint Conference on Measurement and Modeling of Computer Systems*, ACM New York, NY, USA, 15-19 June 2009, pp.157-168.

[8] A. Beloglazov and R. Buyya, "Energy Efficient Resource Management in Virtualized Cloud Data Centers," in *Proceedings of 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, 17-20 May 2010, pp. 826-831.

[9] A. Beloglazov and R. Buyya, "Energy Efficient Allocation of Virtual Machines in Cloud Data Centers," in *Proceedings of 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, 17-20 May2010, pp. 577-578.

[10] J. Berral, Goiri, R. Nou, F. Juli, J. Guitart, R. Gavald, J. Torres, "Towards energy-aware scheduling in data centers using machine learning," in *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking,* Passau, Germany, 13-15 April 2010, pp. 215–224.

[11]K. Hess and A. Newman, "Practical virtualization solutions: virtualization from the Trenches," *Prentice Hall 1st edition*, 2009

[12]B. Loganayagi and Dr. S. Sujatha, "Creating Virtual Platform for Cloud Computing," presented in *IEEE International Conference on Computational Intelligence and Computing Research*, 28-29 December 2010, pp. 1-4.

[13]R. Brown, E. Masanet, B. Nordman, B. Tschudi, A. Shehabi, J. Stanley, *et. al.*, "Report to congress on server and data center energy efficiency: Public law 109-431," *Lawrence Berkeley National Laboratory*, 2008.

[14]A. Beloglazov and R. Buyya, "Energy Efficient Resource Management in Virtualized Cloud Data Centers," in *Proceedings of 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, 17-20 May 2010, pp. 826-831.