Energy Efficient Virtual Machine Consolidation using Dynamic Threshold in Cloud Data Centers

T.Sureshkumar[1], Dr.P.Saveetha[2] and Dr.M.Vijayakumar[3]

[1]Assistant Professor/IT, Nandha College of Technology, Tamil Nadu, India. Email: sureshkumartt@gmail.com

[2]Associate Professor/IT, Nandha College of Technology, Tamil Nadu, India. Email: saveepme@gmail.com

[3]Professor/CSE, Nandha College of Technology, Tamil Nadu, India. Email: tovijayakumar@gmail.com

Abstract - Cloud computing offers value oriented IT services to users universally. Based on pay-as-you-go model, it allows hosting of pervasive applications from shopper, scientific, and business domains. However, knowledge centers hosting Cloud applications consume vast amounts of power, contributory to high operational prices and carbon footprints to the surroundings. Therefore, we want inexperienced Cloud computing solutions that may not solely minimize operational prices however conjointly cut back the environmental impact. Virtualization allows Cloud providers to deal with the energy ineffectiveness by creating multiple Virtual Machine (VMs) instances on a physical server, thus improving the utilization of resources and increasing the Return on Investment (ROI). The reduction in energy consumption can be achieved by switching idle nodes off, thus eliminating the idle power consumption. Moreover, by using live migration the VMs can be dynamically consolidated on the minimal number of physical nodes according to their current resource requirements. However, economical resource management in Clouds is not trivial, as fashionable service applications usually expertise extremely variable workloads inflicting dynamic resource usage patterns. Therefore, aggressive consolidation of VMs will cause performance degradation in application encounters increasing demand leading to an increase of resource usage. Also, Cloud providers have to deal with the energy- performance tradeoff. The focus of this work is on energy and performance efficient resource management strategies that can be applied in a virtualized Cloud data center by a provider.

Keywords - Virtual Machine, Return on Investment, Data Centers, Cloud, FB-DIMM, LLNL, ICT, Quality of Service.

I. Introduction

Cloud computing is a new phase in computing that allows users to access storage or computing power from geographically distant location using Internet connection. According to National institute of Standards and technology (NIST 2011) the definition of Cloud Computing is as follows: “Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage , applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction”.

This cloud model consists of 5 essential characteristics, 3 service models, and 4 deployment models. Cloud computing is maybe the foremost value economical technique to use, maintain and upgrade. Traditional desktop software system prices firms tons in terms of finance. Adding up the licensing fees for multiple users will persuade be terribly high-ticket for the institution involved. The cloud, on the opposite hand, is available at much cheaper rates and hence, can significantly lower the company’s IT expenses. Besides, there are many one-time-payments, pay-as-you-go and other scalable options available, which make it very reasonable for the company in question.

Power and Energy Models: To understand power and energy management mechanisms, it is essential to clarify the terminologies. Electric current is that the flow of electrical charge measured in amperes. Amperes outline the quantity of electrical charge transferred by a circuit per second. Power and energy will be outlined in terms of labor that a system performs. Power is that the rate at that the system performs the work, while energy is the total amount of work performed over a period of time. Power and energy square measure measured in watts (W) and heat unit (Wh), respectively. Work is done at the rate of 1W when 1A is transferred through a potential difference of 1V. A kilowatt-hour (kWh) is that the quantity of energy akin to an influence of one power unit (1000 W) being applied for one hour. Formally, power and energy will be outlined as in (1) and (2):

$$P=W/T \quad (1)$$

$$E= PT \quad (2)$$

where P is power, T is a period of time, W is the total work performed during that period of time, and E is energy. The difference between power and energy is very important because a reduction of the power consumption does not always reduce the consumed energy. For example, the power consumption can be decreased by lowering the CPU performance. However, during this case, a program may require longer time to complete its execution consuming the same amount of energy. On one hand, a reduction of the peak power consumption results in decreased costs of the infrastructure provisioning, such as costs associated with capacities of UPS, PDU, power generators, cooling system, and power distribution equipment. On the other hand, decreased energy consumption leads to a reduction of the electricity bills. The energy consumption will be reduced briefly by dynamic power management (DPM) techniques or by applying static power management (SPM).

DPM utilizes the data of the $64000 time resource usage and application workloads to optimize the energy consumption. However, it doesn't essentially decrease the height power consumption. In distinction, SPM includes the usage of highly efficient hardware equipment, such as CPUs, disk storage, network devices, UPS, and power supplies. These structural changes typically cut back each the energy and peak power consumption.

Static and Dynamic Power Consumption: The main power consumption in complementary metal-oxide-semiconductor (CMOS) circuits comprises static and dynamic power. The static power consumption, or leakage power, is caused by leakage currents that are present in any active circuit, independently of clock rates and usage scenarios. This static power is especially determined by the kind of transistors use for the technology. The reduction of the static power needs enhancements of the low-level system style. Dynamic power consumption is created by circuit activity (i.e., transistor switches, changes of values in registers, etc.) and depends mainly on a specific usage scenario, clock rates, and I/O activity. The sources of the dynamic power consumption are short-circuiting current and switched capacitance. Short-circuit current causes solely 10–15% of the full power consumption and then way no means have been found to scale back this worth while not compromising the performance. Switched capacitance is that the primary supply of the dynamic power consumption; so, the dynamic power consumption will be outlined as in (3):

$$P_{dynamic} = aCV^2f \quad (3)$$

where ‘a’ is that the change activity, C is that the physical capacitance, V is the voltage and f is that the clock frequency. The values of change activity and capacitance are determined by the low-level system style. The combined reduction of the supply voltage and clock frequency lies in the roots of the widely adopted DPM technique called dynamic voltage and frequency scaling (DVFS). The main plan of this method is to designedly down-market the central processing unit performance, once it's not totally utilized, by decreasing the voltage and frequency of the CPU that in the ideal case should result in a cubic reduction of the dynamic power consumption. DVFS is supported by hottest CPUs together with mobile, desktop, and server systems.

Sources of Power Consumption: According to data provided by Intel Labs, the main part of power consumed by a server is accounted for the CPU, followed by the memory and losses due to the power supply inefficiency. The data show that the central processing unit not dominates power consumption by a server. This resulted from the continual improvement of the central processing unit power potency and application of power-saving techniques (e.g., DVFS) that change active low-power modes. In these modes, a central processing unit consumes a fraction of the whole power, while preserving the ability to execute programs. As a result, current desktop and server CPUs can consume less than 30% of their peak power in low activity modes, leading to dynamic power range of more than 70% of the peak power. In contrast, dynamic power ranges of all other server’s components are much narrower: less than 50% for dynamic random access memory (DRAM), 25% for disk drives, 15% for network switches, and negligible for other components. The reason is that only the CPU supports active low-power modes, whereas other components can only be completely or partially switched off. However, the performance overhead of a transition between active and inactive modes is substantial.

For example, a disc drive during a spun-down, deep-sleep mode consumes almost no power, but a transition to active mode incurs a latency that is 1000 times higher than the regular access latency. Power unskillfulness of the server’s elements within the idle state results in a slim overall dynamic power varies of half-hour. This means that if a server is totally idle, it will still consume over seventieth of its peak power. Another reason for the reduction of the fraction of power consumed by the CPU relatively to the whole system is the adoption of multi-core architectures. Multi-core processors are much more efficient than conventional single-core processors. For example, servers built with recent Quad-core Intel Xeon processor can deliver 1.8 teraflops at the peak performance, using less than 10 kW of power.

To compare with, Pentium processors in 1998 would consume about 800 kW to achieve the same performance. The adoption of multi- core CPUs together with the increasing use of virtualization technologies and data-intensive applications resulted within the growing quantity of memory in servers.

In contrast to the CPU, DRAM has a narrower dynamic power range and power consumption by memory chips is increasing. Power consumption by server’s components Memory is packaged in dual in-line memory modules (DIMMs), and power consumption by these modules varies from 5 to 21 W per DIMM, for DDR3 and fully buffered DIMM (FB-DIMM) memory technologies. Power consumption by a server with eight 1GB DIMMs is concerning 80 W.

Modern massive servers presently use thirty two or sixty four DIMMs that results in power consumption by heart above by CPUs. Most of the ability management techniques area unit centered on the central processing unit but, the constantly increasing frequency and capacity of memory chips raise the cooling requirements apart from the problem of high energy consumption. These facts make memory one of the most important server components that have to be efficiently managed. New techniques and approaches to the reduction of the memory power consumption have to be developed in order to address this problem. Power supplies transform alternating current (AC) into direct current (DC) to feed server’s components. This transformation leads to significant power losses due to the inefficiency of the current technology. The efficiency of power supplies depends on their load. They achieve the highest efficiency at loads within the range of 50–75%. However, most information centers unremarkably produce a load of 10–15% wasting the bulk of the consumed electricity and resulting in the common power losses of 60–80%.As a result, power supplies consume at least 2% of the US electricity production. More economical power provider will save over a 1/2 the energy consumption. The problem of the low average utilization additionally applies to disk storages, particularly once disks area unit connected to servers during an information center.

However, this can be addressed by moving the disks to an external centralized storage array. Nevertheless, intelligent policies have to be used to efficiently manage a storage system containing thousands of disks. This creates another direction for the research work aimed at the optimization of the resource, power, and energy usage in server farms and data centers.

II. RELATED WORKS AND PROBLEM ANALYSIS

In the work [1], Author proposed a new technique called minimization of power consumption in a heterogeneous cluster of computing nodes serving multiple web-applications. The main technique applied to reduce power consumption is concentrating the employment to the minimum of physical nodes and switch idle nodes off. This approach needs handling the power/performance trade-off, as performance of applications can be degraded due to the workload consolidation. Requirements to the throughput and execution time of applications are defined in SLAs to ensure reliable QoS. The projected formula sporadically monitors the load of resources (CPU, disk storage and network interface) and makes decisions on switching nodes on/off to minimize the overall power consumption, while providing the expected performance. The actual load leveling isn't handled by the system and has got to be managed by the applications. The formula runs on a master node, which creates a Single Point of Failure (SPF) and may become a performance bottleneck in a large system. In addition, the authors have pointed out that the reconfiguration operations are time-consuming, and the algorithm adds or removes only one node at a time, which may also be a reason for slow reaction in large-scale environments. The proposed approach can be applied to multi-application mixed-workload environments with fixed SLAs.

In the work [2], Author proposed a new technique, called power-efficient allocation of VMs in virtualized heterogeneous computing environments. They have leveraged the min, max and shares parameters of VMM, which represent minimum, maximum and proportion of the CPU allocated to VMs sharing the same resource. Similarly to the approach suits only enterprise environments as it does not support strict SLAs and requires the knowledge of application priorities to define the shares parameter. Other limitations are that the allocation of VMs is not adapted at run-time (the allocation is static) and no other resources except for the CPU are considered during the VM reallocation.

In the work [3], Author proposed a new technique, called scheduling for multi-tiered web-applications in virtualized heterogeneous systems to minimize energy consumption, while meeting performance requirements. The authors have investigated the result of performance degradation thanks to high utilization of various resources once the employment is consolidated. They have found that the energy consumption per dealings ends up in a ‘U’-shaped curve, and it is possible to determine the optimal utilization point. To handle the improvement over multiple resources, the authors have proposed a heuristic for the multidimensional packing problem as an algorithm for the workload consolidation. However, the projected approach is employment kind and application dependent, whereas our algorithms are independent of the workload type, and thus are suitable for a generic Cloud environment.

III. METHODOLOGY

The System operation can be divided into two parts: (1) selection of VMs that have to be migrated to optimize the allocation; and (2) placement of the VMs selected for migration and new VMs required by the users on physical nodes. Discuss these parts in the following sections.

A. VM Selection

Fixed Utilization Thresholds: In the previous work it has proposed four heuristics for choosing VMs to migrate. The first heuristic, Single Threshold (ST), is based on the idea of setting an upper utilization threshold for hosts and placing VMs while keeping the total utilization of the CPU below this threshold. The aim is to preserve free resources to stop SLA violation thanks to consolidation in cases once the resource demand by VMs will increase.

At each time frame all the VMs are reallocated using the Modified Best Fit Decreasing (MBFD) algorithm with an 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 the total utilization of the CPU by all the VMs between these thresholds. If the CPU utilization of 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 the idle power consumption. If the utilization exceeds the upper threshold, some VMs have to be migrated from the host to reduce the utilization in order to prevent potential SLA violation. Three policies are proposed for choosing VMs that have to be migrated from an over-utilized host.

- Minimization of Migrations (MM) – migrate the least number of VMs to minimize migration overhead.
- Highest Potential Growth (HPG) – migrate VMs that have rock bottom usage of central processing unit comparatively to request so as to reduce total potential increase of the use and SLA violation.

- Random Choice (RC) – choose the necessary number of VMs randomly.

Dynamic Utilization Thresholds: As mentioned earlier, fixed values for the thresholds are unsuitable for an environment with dynamic and unpredictable workloads, in which different types of applications can share a physical resource. The system should be able to automatically adjust its behavior depending on the work- load patterns exhibited by the applications. Therefore, it propose a novel technique for auto-adjustment of the utilization thresholds based on a statistical analysis of the historical data collected during the lifetime of VMs. First of all, assume that the CPU utilization created by each VM can be described by a random variable (uj) with a particular distribution, which persists at least over some recent period of time. In this case, the CPU utilization of a host can be represented by a random variable (Ui), which is a sum of utilizations by m number of VMs allocated to this host. Then assume that as the distributions created by different VMs are different, the distribution of the host's utilization is approximately normal and can be modeled by the t-distribution. We cannot predict the CPU utilization of a physical node in the future; however, we can calculate characteristics of the distribution over some recent period of time, such as the sample mean (Ui) and standard deviation (sUi).

The advantage of collecting the data for each VM separately and then using the summation is that a VM is mi- grated together with the data of its resource usage and the data will be actually even after a VM migration. Using this information and the inverse cumulative probability function for the t-distribution (tinvn (P)), it is possible to find out an interval of the CPU utilization, which be reduced with a low probability (e.g. 5%). It can set the upper utilization threshold (Tui) for each host i preserving this amount of spare CPU capacity defined by the lower (Pul) and up- per (Puu) limits of the probability interval as shown in below, where n is the number of data points collected, and $n - 1$ represents the degrees of freedom for the t-distribution. The lower threshold is calculated in a similar way. However, the difference is that a single value is obtained for all the hosts in the system. The idea is to determine the hosts that have lower utilizations relatively to the average value across all the nodes. To tackle the case when all the hosts have low CPU utilizations, it introduces a limit (Ul) to cap the decrease of the lower utilization threshold. In our previous work we have found the value Ul = 30% to be effective for the lower threshold. The reallocation algorithm using the dynamic thresholds (DT) is presented in Algorithm. For the DT algorithm it apply the MM policy for VM selection, as in the previous work it has shown the superiority over the alternatives. The complexity of the algorithm is proportional to the sum of the number of non over-utilized host plus the product of the number of over-utilized hosts and the number of VMs allocation to these over-utilized hosts.

B. VM Placement

The VM placement can be seen as a bin packing problem with variable bin sizes and prices, where bins represent the physical nodes items are the VMs that have to be allocated bin sizes are the available CPU capacities of the nodes and prices correspond to the power consumption by the nodes. As the bin packing problem is NP-hard, to solve it apply a modification of the Best Fit Decreasing (BFD) algorithm that is shown to use no more than 11/9 * OP T +1 bins (where OP T is the number of bins provided by the optimal solution) [15]. In the modification (MBFD) it sort all the VMs in the decreasing order of current CPU utilizations and allocate each VM to a host that provides the least increase of the power consumption caused by the allocation. This allows the leveraging the nodes heterogeneity by choosing the most power-efficient ones first. The pseudo-code for the algorithm is presented in Algorithm. The complexity of the algorithm is n * m, where n is the number of nodes and m is the number of VMs that have to be allocated.

C. Modified Best Fit Decreasing (MBFD)

Input: hostList, vmList Output: allocation of VMs vmList.sortDecreasingUtilization()
foreach vm in vmList do minpower ← MAX allocationHost ← NULL foreach host in hostList do
if host has enough resource for vm then
power ← estimatePower (host, vm)
if power < manpower then
allocatedHost ← host manpower ← power
if allocationHost ≠ NULL then allocate vm to allocatedHost return allocation

D. Dynamic Thresholds (DT)

Input: hostList, vmList Output: migrationList vmList.sortDecreasingUtilization()
foreach h in hostList do hUtil ← h.util() bestFitUtil ← MAX
while hUtil > h.upThresh() do foreach vm in vmList do
if vm.util() > hUtil – h.upThresh() then
t←vm.util() – hUtil + h.upThresh()
if t < bestFitUtil then bestFitUtil ← t bestFitVm ← vm
else

if bestFitUtil = MAX then
bestFitVm ← vm
break
hUtil ← hUtil – bestFitVm.util() migrationList.add(bestFitVm) vmList.remove(vm)
if hUtil < lowThresh () then migrationList.add(h.getVmlist()) vmlist.remove(h.getVmList()) return migrationList

IV. EXPERIMENTAL STUDY

As the target system is a generic Cloud computing environment, it is essential to evaluate it on a large-scale virtualized data center infrastructure. However, it is difficult to conduct large-scale experiments on a real infrastructure, especially when it is necessary to reproduce the experiment with the same conditions to compare different algorithms. Therefore, a simulation has been chosen as the way to judge the planned algorithms.

The CloudSim toolkit 2.0 has been chosen as a simulation platform, as it is a modern simulation framework aimed at Cloud computing environments. In contrast to alternative simulation toolkits (e.g. SimGrid, GangSim), it allows the modeling of virtualized environments, supporting on-demand resource provisioning, and their management. It has been extended so as to modify power-aware simulations and dynamic workloads, as the core framework does not provide these capabilities. The enforced extensions are enclosed within the 2.0 version of the CloudSim toolkit.

V. CONCLUSION AND FUTURE ENHANCEMENT

In this paper, a model is proposed for the data center with the given requirements by CloudSim to apply energy efficient resource management strategies such as dynamic consolidation of virtual machines and switching off the idle node. The Dynamic Threshold (DT) and Modified Best Fit Decreasing (MBFD) Algorithms are used to reduce the energy consumption by the nodes in cloud data centers. The experimental results show that the proposed technique outperforms SLA violation and can be limited to < 1% and number of virtual machine migrations, while providing a similar level of energy consumption. An efficient algorithm is identified and studied for the energy efficient virtual machine consolidation. For the future work, consideration of multiple system resources, such as memory and network interface to evaluate the proposed system in the real Cloud infrastructure.

REFERENCES

[1] A.Beloglazov and R. Buyya. Energy efficient allocation of virtual machines in cloud data centers. In Proc. of the 10th IEEE/ACM Intl. Symp. on Cluster, Cloud and Grid Computing (CCGrid 2010), 2010.
[2] R. N. Calheiros. CloudSim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. Software: Practice and Experience, Wiley Press, NY, USA, 2010.
[3] VMware Inc. VMware distributed power management concepts and use, 2010.
[4] Y. Song. Multi-Tiered On-Demand resource scheduling for VM- Based data center. In Proc. of the 2009 9th IEEE/ACM Intl. Symp. On Cluster Computing, pp:148-155, 2009.
[5] C. Vecchiola and Aneka: a software platform for .NET-based cloud computing. High Performance & Large Scale Comp., Advances in Parallel Computing,pp 267-295, 2009.
[6] W.Voorsluys. Cost of virtual machine live migration in clouds: A performance evaluation. In Proc. of the 1st Intl. Conf. on Cloud Computing, pp: 254-265, 2009.
[7] V.S. Sureshkumar, A.Chandrasekar, "Fuzzy-GA Optimized Multi-Cloud Multi-Task Scheduler For Cloud Storage And Service Applications" International Journal of Scientific & Engineering Research, vol. 4, no. 3, pp. 1-7, March 2013
[8] D. Kusic. Power and performance management of virtualized computing environments via lookahead control. Cluster Computing, vol. 12. no. 1, pp. 1-15, 2009.
[9] A. Verma and pMapper: power and migration cost aware application placement in virtualized systems. In Proc. of the 9th ACM/IFIP/USENIX Intl. Conf. on Middleware, pp. 243-264, 2008.
[10] R. Buyya. Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities. In Proc. of the 10th IEEE Intl. Conf. on High Performance Computing and Communications (HPCC'08), 2008.
[11] R. Brown. Report to congress on server and data center energy efficiency: Public law 109-431. Lawrence Berkeley National Laboratory, 2008.
[12] R.Nathuji and K. Schwan. Virtualpower: Coordinated power management in virtualized enterprise systems. ACM SIGOPS Operating Systems Review, vol. 41, no. 6, pp. 265-278, 2007.
[13] V.S. Suresh kumar “Extended Framework For Dynamic Resource Allocation Using Asjs Algorithm In Cloud Computing Environment", International Journal on Engineering Technology and Sciences vol.1, no. 8, pp. 1-7, 2014.
[14] X.Fan. Power provisioning for a warehouse-sized computer, In Proc. of the 34th Annual Intl. Symp. On Computer Architecture, pp: 13-23, 2007.
[15] 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, vol. 7, no. 4, pp. 321-331, 1991.
[16] Nandagopal S, Karthik S, and Arunachalam VP., “Mining of meteorological data using modified apriori algorithm”, European Journal of Scientific Research, vol. 47, no. 2, pp.295-308, 2010.
[17] Saveetha P and Arumugam S, “Study on Improvement in RSA Algorithm and its Implementation”, International Journal of Computer & Communication Technology, Vol.3 No.6, pp:78, 2012.
[18] Nithya K, KalaivaaniPCD and ThangarajanR, “An enhanced data mining model for text classification”, International Conference on Computing, Communication and Applications, pp:1-4,2012.
[19] Vijayakumar, M., Prakash, S. and Parvathi, R.M.S. “Inter Cluster Distance Management Model with Optimal Centroid Estimation for K-Means Clustering Algorithm, “WSEAS Transactions on Communications, Issue 6, Vol. 10, pp. 182-191, June 2011.
[20] Vijayakumar M and Parvathi RMS, “Concept mining of high volume data streams in network traffic using hierarchical clustering”, European Journal of Scientific Research, Vol.39,No.2, pp:234-242,January 2010.
[21] Prakash S, Vijayakumar M and Parvathi RMS, “A novel method of mining association rule with multilevel concept hierarchy”,International Journal Computer Application(IJCA), pp::26-29,2011.

[22] Nandagopal S, Arunachalam VP and Karthik"A Novel Approach for Mining Inter-Transaction Item sets", European Scientific Journal,Vol.8, No:14, pp::92-108,2012.

[23] Prakash S and Vijayakumar M, "An effective network traffic data control using improved Apriori rule mining", Circuits and Systems, vol. 7, no. 10, pp. 3162-3173, August 2016.

[24] Gokulraj P and Kiruthikadevi K, "Revocation and security based ownership deduplication of convergent key creating in cloud", International Journal of Innovative Research in Science, Engineering and technology. Vol. 3, Issue 10, ISSN: 2319-8753, October 2014.

[25] Saveetha P, Arumugam S and Kiruthikadevi K, "Cryptography and the Optimization Heuristics Techniques", Int. Journal of Advanced Research in Computer Science and Software Engg , volume. 4, Issue.10, ISSN: 2277 128X, October 2014.

[26] Vijayakumar M and Prakash.S"An Improved Sensitive Association Rule Mining using Fuzzy Partition Algorithm", Asian Journal of Research in Social Sciences and Humanities,Vol.6,No.6, pp:.969-981,2016.

[27] Prakash S and Vijayakumar M, "Risk Assessment in Cancer Treatment using Association Rule Mining Techniques", Asian Journal of Research in Social Sciences and Humanities, Vol.6,No.10,2016.

[28] Preethi, B.C. and Vijayakumar, M. " A Novel Cloud Integration Algorithm(CIA) for Energy Efficient High Performance Computing Applications in Big Data Multimedia Applications", Romanian Journal of Information Science and Technology, vol. 2, no.1, pp. 1-11, March 2018.